ARL-SR-0404 ● SEP 2018



US Army Research Laboratory

Micro Autonomous Systems and Technology (MAST) Final Annual Report for Program Capstone

by Brett Piekarski, Allison Mathis, William Nothwang, David Baran, Christopher Kroninger, Brian Sadler, Larry Matthies, Vijay Kumar, Inderjit Chopra, J Sean Humbert, Kamal Sarabandi, and Susan Wright

Approved for public release; distribution is unlimited.

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

ARL-SR-0404 ● SEP 2018



US Army Research Laboratory

Micro Autonomous Systems and Technology (MAST) Final Annual Report for Program Capstone

by Brett Piekarski, Christopher Kroninger, and Brian Sadler, ***Vehicle Technology Directorate, ARL***

Allison Mathis and William Nothwang, ***Sensors and Electron Devices Directorate, ARL***

David Baran, ***Computational and Information Sciences Directorate, ARL***

Larry Matthies, ***NASA Jet Propulsion Laboratory, Pasadena, CA***

Vijay Kumar, ***University of Pennsylvania, Philadelphia, PA***

Inderjit Chopra, ***University of Maryland, College Park, MD***

J Sean Humbert, ***University of Colorado, Boulder, CO***

Kamal Sarabandi, ***University of Michigan, Ann Arbor, MI***

Susan Wright, ***BAE Systems, Nashua, NH***

Approved for public release; distribution is unlimited.

REPORT DOCUMENTATION PAGE	Form Approved OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED (From - To)
September 2018	Special Report	2008–2018

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Micro Autonomous Systems and Technology (MAST) Final Annual Report for Program Capstone	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)	5d. PROJECT NUMBER
Brett Piekarski, Allison Mathis, William Nothwang, David Baran, Christopher Kroninger, Brian Sadler, Larry Matthies, Vijay Kumar, Inderjit Chopra, Sean Humbert, Kamal Sarabandi, and Susan Wright	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
US Army Research Laboratory ATTN: RDRL-SER-L 2800 Powder Mill Road Adelphi, MD 20783-1138	ARL-SR-0404

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The vision of the Micro Autonomous Systems and Technology (MAST) program was to enhance tactical situational awareness in urban and complex terrain by enabling the autonomous operation of a collaborative ensemble of multifunctional, mobile microsystems through collaborative and multidisciplinary research. Through a comprehensive basic research program, the MAST consortium developed innovative technologies to enable Soldier-portable, scaled unmanned aerial systems (UASs) and unmanned ground vehicles (UGVs). Specific directions of advancement included scaling down autonomous capabilities to handheld-sized platforms; advancing GPS-denied mapping and navigation methods; driving the development of vision-based approaches for state estimation, control, and navigation; increasing the operational speed and aggressive control for small UASs; developing methods for collaborative behaviors between UAS and UGV autonomous systems; enabling bio-inspired behaviors such as perching and complex terrain agility; developing and validating physics-based computational fluid dynamics simulation tools for MAST-scaled and highly maneuverable UASs; and advancing new technologies to address the size, weight, power, and algorithm efficiency of sensing and communications technologies.

15. SUBJECT TERMS

Micro Autonomous Systems and Technology, MAST, situational awareness, mobile platforms, autonomous systems

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Allison Mathis
Unclassified	Unclassified	Unclassified	UU	470	19b. TELEPHONE NUMBER (Include area code) 301-394-5518

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

Contents

1. Micro Autonomous Systems and Technology (MAST) Collaborative Technology Alliance (CTA) Vision 1

1.1 Driving Scenarios 1

1.2 Characteristics of a MAST System 3

1.3 Technical Challenge 3

1.4 MAST CTA Program – Structure and Organization and Centers/Thrusts Areas 4

1.5 Scientific Advancements Under the MAST CTA 6

1.6 Impact of Future Systems 8

1.7 References 9

2. Mobility, Control, and Energetics Thrust 10

2.1 Vision of the Mobility, Control, and Energetics Thrust 10

2.2 Technical Challenges 10

2.3 Background 11

2.4 Thrust Organization 13

2.5 Technical Program 13

2.5.1 Aerial Mobility 14

2.5.2 Ground Mobility 15

2.5.3 Actuation and Power 16

2.5.4 Dynamics and Control 17

2.6 Significant Accomplishments 18

2.7 Metrics of Success 20

2.7.1 Some High-Impact Publications 20

2.8 Challenges for the Future 20

2.9 MCE Task Summary 22

2.10 References 22

MCE Research Task Reports 24

3. Sensing, Perception, and Processing (SPP) Thrust 138

Approved for public release; distribution is unlimited.

3.1 Vision for Sensing, Perception, and Processing 138
3.2 Technical Challenges 138
3.3 Thrust Organization 139
3.4 Technical Challenges and Approach 140
3.4.1 Sensing and Estimation 140
3.4.2 3-D Perception and Mapping 141
3.4.3 ISR 144
3.5 Significant Accomplishments 145
3.6 Challenges for the Future 146
3.7 SPP Task Summary 147
3.8 References 147

SPP Research Task Reports 150

4. Communication, Networking, and Coordination (CNC) Thrust 217
4.1 Research Vision 217
4.2 Army Relevance and Technical Potential 218
4.3 Technical Challenges 218
4.3.1 Novel Radios 219
4.3.2 Communication and Control Software 220
4.3.3 Group Behaviors 221
4.4 Technical Plan 222
4.4.1 Novel Radios 222
4.4.2 Communication and Control Software 223
4.4.3 Group Behaviors 224
4.5 Key Program Technical Accomplishments 225
4.6 Challenges for the Future 227
4.6.1 Heterogeneity 227
4.6.2 Resiliency 228
4.6.3 Learning 228
4.6.4 High-Level Task Specification 229
4.6.5 Autonomous Networking 229
4.7 CNC Task Summary 230
4.8 Section Bibliography 230

CNC Research Task Reports **232**

5. Joint Experiments (JX) Thrust **311**

5.1 Introduction 311
5.1.1 Research Vision 311
5.1.2 Army Relevance and Technical Potential 312
5.2 Technical Challenges: 2008–Present 312
5.2.1 Stage 1: FY10, Pairwise Component Integration Tasks 312
5.2.2 Stage 2: FY11–12, JX Thrust Established with 4 Focus Areas Leading to Year 5 Capstone 314
5.2.3 Stage 3: FY13–14, Option Years with New Tasks Selected from White Papers 316
5.2.4 Stage 4: FY15–17, Capability Maturation 318
5.3 Technical Plans and Progress in the FY15–17 Phase 320
5.3.1 High-Speed Obstacle Avoidance 320
5.3.2 Autonomous MAV Perch-and-Stare 323
5.3.3 Mapping with Heterogeneous Robot Teams 325
5.3.4 Airborne Sense-and-Avoid for MAVs 327
5.3.5 Autonomous Navigation on MAST-Scale Aircraft 330
5.4 Challenges for the Future 332
5.5 JX Task Summary 333
5.6 References 333

JX Research Task Reports **336**

6. MAST Program Bibliography **358**

List of Symbols, Abbreviations, and Acronyms **460**

Distribution List **463**

1. Micro Autonomous Systems and Technology (MAST) Collaborative Technology Alliance (CTA) Vision

The Micro Autonomous Systems and Technology (MAST) Collaborative Technology Alliance (CTA) was initiated in 2008 by the US Army Research Laboratory (ARL) to explore and develop the underpinning science to realize future concepts in collaborative and Soldier-portable micro autonomous systems for enhanced tactical situational awareness in urban and complex terrain. The operational vision for MAST focuses on autonomous operation over the last 100 m from the Soldier or squad to, into, and around a non-line-of-sight (NLOS) complex environment to provide the dismounted Soldier extended-reach intelligence, surveillance, and reconnaissance (ISR) operations. This may encompass GPS-denied outdoor autonomous travel through extreme environments (e.g., wind gusts) to an indoor complex environment (for physical movement, communication, and perception) such as a building, cave, or analogous aboveground or underground structure. Another driver is the operational tempo at which these MAST systems can operate with the goal of maintaining Soldier operational tempo even in complex environments or being able to move with the Soldier or squad as they move, potentially creating a mobile ISR bubble around them. These operations may take systems indoors, outdoors, or underground and through both structured and unstructured environments.

For the scope of the research, the total mission duration for approaching and performing ISR tasks at an objective is considered to be less than an hour; for mobile ISR operations that move with the Soldier or squad as it moves through the environment, the total mission duration is considered to be up to 72 h and beyond. This vision drives the major research themes of the program without introducing difficulties that are out of reach or irrelevant to the research.

1.1 Driving Scenarios

This vision was further defined by the establishment of four driving scenarios to help shape the research focus within the MAST CTA. These scenarios are designed to address the vision of operation in realistic, real-world environments, progressing from relatively benign conditions in Scenario 1 to far more challenging conditions and operational tempo in Scenario 4. Notional expansions of these scenarios are given briefly in what follows. The base period of the program focused primarily on Scenario 1, plus some enabling elements of Scenario 2. Additional challenges presented by Scenarios 3 and 4 have been increasingly addressed over the latter half of the program.

Scenario 1 *(Small-unit building search): A small number of platforms start from approximately 100 m away from a small building with nominal daytime lighting conditions, smooth ground surfaces, and quiescent atmosphere*. For example, a heterogeneous mix of air and ground platforms approach the target structure, and some air platforms may carry smaller, deployable ground platforms. Platforms search for and identify an ingress point with optional user assistance (window, door, vent, and hole), then enter and distribute themselves throughout the building to conduct reconnaissance and mapping, and maintain communication connectivity with all units. When the team leaves the building, some units may remain behind, in hidden locations, to conduct ongoing surveillance of the building interior.

Scenario 2 *(Small-unit cave search): Similar to Scenario 1, but with increasing environmental complexity and uncertainty*. The cave or multistory collapsed building has a complex 3-D structure, minimal or no lighting, uneven terrain, some rubble, branching passages, varying ceiling height, and irregular openings for ingress/egress. Some areas are impassable to ground vehicles. There are mild to moderate wind gusts and atmospheric perturbations, especially at ingress and egress points. A heterogeneous team of air and ground vehicles approaches and enters the structure. Once inside, the team traverses the terrain, avoids obstacles, searches, maps, localizes, communicates, and partitions the effort among vehicle types as needed to accomplish the mission under these conditions.

Scenario 3 *(Small-unit perimeter defense): Autonomously plan and execute reconnaissance missions with multiple air and ground platforms over "large" 3-D areas including outdoor and indoor segments (e.g., multiple, multi-floor buildings).* Objectives: 1) identify the extent of buildings or structures along with potential ingress and egress points; 2) maintain communication and position knowledge between all assets, without line of sight (LOS), including inside and outside buildings and underground, with platforms in any orientation and at any height; 3) conduct all-terrain ground mobility in this area, horizontal, vertical, smooth, and rough surfaces; 4) conduct full 3-D air mobility and station keeping in the presence of strong wind gusts, including situations such as flying around corners or over rooftops while coping with varying gusts in those locations; and 5) cooperatively detect, track, and follow people and vehicles moving in this area using visual and/or other sensors.

Scenario 4 *(Soldier- and squad-level mobile ISR): Autonomously plan and execute reconnaissance missions with multiple air and ground platforms at the operational tempo of the Soldier or squad to provide situational ISR as they move through complex 3-D areas, including outdoor and indoor segments (e.g., multiple, multi-floor buildings).* This scenario is an extension of Scenario 3 with the added

requirement that the heterogeneous team of robots can perform the same objectives while moving with the Soldier or squad at their operational tempo through complex environments to maintain a 100-m ISR bubble throughout the duration of the mission. This requires additional technologies and approaches to incorporate heterogeneous teams, and assess risk and increase operational tempo and mission duration.

1.2 Characteristics of a MAST System

To help further shape the research and vision for the program, a few key or desired characteristics of the systems and their capability were developed. The MAST consortium is working toward creating systems of heterogeneous autonomous platforms equipped with miniature sensors to quickly, quietly, and reliably explore and find targets in the scenarios defined by the driving scenarios. To meet the vision of the program, these platforms will exploit size (tradeoffs between payload, minimum exploitable opening, speed, range, and duration), diversity of locomotion (flight and ambulation), multiple units (redundancy, communication network, observation angles), and diversity of sensing (passive and active, local and long range). Therefore, the key characteristics sought within the program are small scale (handheld size or smaller) to enable Soldier-portable systems; single platform autonomy in GPS-denied and complex environments to reduce the need for human control and interactions; and collective behaviors to increase coverage, reduce mission times, and enable behaviors not realizable with individual systems.

1.3 Technical Challenge

The specific tasks within the program arose to address the many technical challenges required to achieve the MAST vision. The initial challenge is to realize aggressive autonomous and operational behaviors in a Solider-portable platform and thus highly constrained in size, weight, power, and processing systems. From there, MAST has moved to investigations into a variety of topics, including robust operation across domains (urban, subterranean, indoor/outdoor, etc.); dynamic and potentially extreme environmental conditions (such as varying lighting, mobile agents, wind gusts, etc.); high-tempo operation through cluttered 3-D environments; localization and operation in austere environments; long-duration GPS-denied navigation; discovery of novel bio-inspired mobility in complex terrain and bio-inspired behaviors like perching; collaborative behaviors across air and ground domains; and communication in complex RF environments.

Unlike most commercial endeavors, the military cannot rely on existing infrastructure. The Department of Defense (DOD) is forced to assume a hostile and

peer-contested, GPS- and communications-denied space with no a priori knowledge of the environment. Military systems are expected to work in extreme environments, some as sparse and featureless as a snow field, desert sand, or a moonless night.

In addition to the previously mentioned challenges, MAST systems are expected to perform military-relevant operations, such as ISR, building search and clearance, route clearance, resupply, perimeter surveillance, and so on. These challenge are defined further within the specific research thrust area overviews and help to define the relevancy of the tasks to the MAST CTA vision and the needs of the Army.

1.4 MAST CTA Program – Structure and Organization and Centers/Thrusts Areas

The MAST CTA was awarded to a consortium of industry and academic partners. Together, the consortium and ARL make up the CTA to jointly perform the underpinning research necessary to realize the MAST CTA vision. The CTA is organized into four research centers, led by four principal members of the consortium. The first three technical centers are focused on *Microsystem Mechanics*, *Processing for Autonomous Operation*, and *Microelectronics*. These centers are led by academic institutions and are referred to as centers. The fourth area is *Integration* and is led by industry. The consortium is augmented by subawardee organizations and principal investigators (PIs) to help round out the innovative research topics within each center.

The principal member for Microsystem Mechanics is the University of Maryland. This research is developing a fundamental understanding of mechanics for small, unmanned air and ground vehicles as needed to obtain desired mobility objectives. The aeromechanics discipline focuses on key elements of microsystem flight; the ambulation discipline focuses on key elements of microsystem ground movement.

The principal member for Processing for Autonomous Operation is The University of Pennsylvania. This center provides the fundamental underpinnings for autonomous operation of distributed, mobile, multimodal sensing microsystems. Areas of investigation include control for mobility and situational awareness, distributed sensing and estimation, algorithmic aspects of communication, and architectures for cooperative reconnaissance by a heterogeneous set of micro air and ground vehicles.

The principal member for Microelectronics is the University of Michigan. The significant constraints on size, weight, and power (SWaP) and processing for MAST-scale platforms require revolutionary approaches to microelectronics for

these systems. Research areas include sensors for navigation and ISR, power systems, radio components, and computational elements.

The Integration Center is focused on system analysis, modeling and experimentation research, and other integration issues. As the lead for the consortium, the principal member for Integration has primary responsibility for articulating and executing a vision for cross-consortium integration. The lead integrates the disparate technologies delivered by the centers for experiments and guides center investigations based on collected data. The principal member for Integration, and the overall lead for the consortium, is BAE Systems. BAE Systems and the Integration Center is supported by the Jet Propulsion Laboratory (JPL).

To develop and manage the overall technical vision for the program, the mission scenarios have been used to derive lists of technical capabilities needed to enable them. These capabilities have been grouped into four cross-cutting thrust areas to help foster collaboration across the centers and with ARL. To further communication and collaboration across the alliance, all of the CTA reviews and documentation, such as this annual report, are organized by thrust area. The four research thrust areas within the MAST CTA and the leading organization for the thrust area are as follows:

- Mobility, Control, and Energetics (MCE) (University of Maryland)
    - MCE develops a fundamental understanding of mechanics for small unmanned air and ground vehicles, as needed, to obtain desired mobility objectives. The aeromechanics discipline focuses on key elements of microsystem flight; the ambulation discipline focuses on key elements of microsystem ground movement.
- Sensing, Perception, and Processing (SPP) (University of Michigan)
    - SPP provides the fundamental underpinning for autonomous operation of distrusted, mobile, multimodal sensing microsystems. Areas of investigation include control for mobility and situational awareness, distributed sensing and estimation, algorithmic aspects of communication, and architectures for cooperative reconnaissance by a heterogeneous set of micro air and ground vehicles.
- Communication, Networking, and Coordination (CNC) (University of Pennsylvania)
    - CNC is concerned with the use of communication for carrying out mission objectives, optimizing the network connectivity among the heterogeneous ensemble of robots using the ratio, and coordinating the

functions of each robot in the ensemble to enable the functioning of the team and optimize the mission objectives.

- Joint Experiments (JX) (JPL)
    - The three core research thrusts (MCE, SPP, and CNC) address capability gaps at the level of component technologies. The JX thrust serves as a focal point to integrate and evaluate technical capabilities on platforms performing elements of mission-like scenarios, through a set of research tasks that involve collaboration between multiple centers and between the consortium and ARL.

The MAST consortium has worked collaboratively to develop a management structure that embraces this paradigm shift through an integrated research plan that incorporates ARL researchers with the consortium and cross-cutting themes to join the four research centers together as one.

A strong collaboration among government, industry, and university researchers is aimed at the following:

1) Guiding the research toward relevant technology development across the consortium.
2) Providing all members with a voice in the management of the consortium and research goals.
3) Providing an environment that promotes unforeseen technology breakthroughs by operating at the nexus of the technology domains not solely within each domain.
4) Enabling free flow of information and expertise among the researchers and development of personal bonds that will stimulate collaboration.
5) Providing a measured and effective path to transition of innovative MAST technologies.
6) Engaging in flexible and active staff rotation and education programs to encourage researchers to work and learn with their peers.

1.5 Scientific Advancements Under the MAST CTA

The MAST CTA has been demonstrating "the art of the possible" since 2008, its mission to inform and enable future Solider-portable autonomous systems. Research within the program has driven significant advancement in a number of different areas. MAST performers have been demonstrating on-platform

autonomous behaviors since the early days of the program, using 2-kg surrogate vehicles until smaller ones could be developed. Advances within the program, along with commercial advancements in processing and memory, have significantly driven down the scale of autonomy-capable unmanned aerial vehicles (UAVs). Performers have demonstrated a 250-g autonomous system

MAST systems have gone from fairly slow-tempo operation, around 2 m/s, in relatively simple, structured 2-D+ environments to operating above 10 m/s in full 3-D cluttered environments. This work contributed to the inspiration for the Defense Advanced Research Projects Agency (DARPA) Fast Lightweight Autonomy (FLA) program, which is stretching this goal even further, to 20 m/s, in complex, 3-D environments.

MAST has produced significant advancements in transitioning between indoor and outdoor environments; in aggressive maneuvers through extreme transitions in orientation and state; in long-duration GPS-denied navigation; in navigation through austere, featureless environments; and in robustness to external environmental conditions, such as wind gust rejection.

At the beginning of the program, technology did not exist to allow platforms to interact with the physical environment around them, and they could not perform advanced biologically inspired behaviors, such as perching. New understanding of biological systems and their control methodologies, and in the materials and approaches required for gripping onto a variety of surfaces, has led to demonstrable vision-only based systems that can perch on surfaces ranging from glass to stucco, on vertical walls, overhead poles, and branches.

Major advances in the validation and sophistication of a physics-based computational fluid dynamics (CFD) simulation tool for MAST-scale flying platforms has lead in an increased understanding of nontraditional approaches for flight, and of the aerodynamics associated with insect and avian flight. It has led to the realization of novel platforms within the program, notably a stable MAST-scaled cyclocopter. Using these tools, the performers were able to scale their original 500-g system down to a fully functional 29-g version.

The MAST program saw significant advancements in the understanding and underpinning theory of collaborative behavior between systems. Early in the program, performers demonstrated 20 MAST-scale systems flying in formation and through obstacles, which has provided the basis for many of the large-scale demonstrations seen throughout the community today. This work was primarily accomplished using external state-estimation infrastructure and centralized processing, on fairly homogenous platforms. Recent work includes collaborative

autonomous navigation and mapping performed by heterogeneous groups of air and ground systems, systems with all processing and control performed onboard, and systems that perform state estimation and relative positioning using only visual perception, without the need for external infrastructure or GPS.

The MAST program has also seen significant discovery and demonstration in terramechanics for mobility of MAST-scale ground systems over a variety of terrains. This work enables transitions between surfaces, increases the efficiency of motion through complex physical environments and is the result of studies of biological systems, such as lizards and insects, to discover the approaches for highly maneuverable systems in complex terrain. It has led to a novel methods and demonstrations of passive, body-shape-based methods for maneuvering through high aspect ratio obstacles (such as grass), self-righting, new gripper technologies, the ability to transition between surfaces, mobility in compressed space and over granular media, the use of appendages and tails for high-speed stability and maneuvers over rugged terrain, and novel hopping platforms that are working toward mimicking the agility of arboreal mammals, which rapidly and repeatedly jump vertical distances of 2 m, starting from a rest stage.

Finally, the program has made significant advances in energy-efficient electronics. These include novel electronics for processing optic flow and stereo imagery needed for high-speed obstacle avoidance and navigation, as well as methods of fuse them together for improved scene perception. There are novel sensors, such as a 5-g scanning radar, which could compliment electro-optic (EO)/IR systems for obstacle detection and navigation in extreme environments. Work includes novel microelectromechanical systems (MEMS) hair-like accelerometer arrays for increased platform state estimation, MEMS hair-like hydraulic sensors for detection and subsequent rejection of wind gusts. In addition, the program has produced methodologies to scale down very high frequency (VHF) antennas into the 40-MHz range to provide a robust multispectral approach to communications and a method of realize communications deep within complex RF environments that are not normally reachable by traditional Wi-Fi or zigbee radios.

1.6 Impact of Future Systems

When the MAST CTA started in 2008, micro autonomous systems were not on any Army or DOD autonomous systems roadmap and no programs of record were in place. Advancements out of the MAST CTA and other programs have contributed significantly to the literature and state of the art in the field of small Soldier-portable systems The challenging initial goal of the program to achieve collaborative and autonomous behaviors in handheld-sized platforms will be realized. These

advancements in the field have influenced the DOD community and one can now find references to micro autonomous systems in documents such as the *Unmanned Systems Integrated Roadmap*,[1] *Unmanned Ground Systems Roadmap,*[2] and the recently released *Army Robotic and Autonomy Systems Strategy,*[3] which includes statements such as "to *increase situational awareness* for dismounted forces in the near term, the Army procures more man-portable RAS at lower echelons, capitalizing on increased endurance, sensory-obstacle avoidance, autonomy and miniaturization for small UGSs and UASs that enable tactical forces to The US Army Robotic and Autonomous Systems Strategy make contact with threats on their own terms" and "squads and platoons equipped with small RAS in urban terrain make contact on their own terms, thus reducing the need for formations to maintain the traditional 6:1 attacker-to-defender ratio commonly associated with conventional urban combat operations." In addition, the November 2013 Army approved Capability Production Document (CPD) for the Family of Small UAS includes a short-range reconnaissance (SRR) small unmanned aircraft system (SUAS). The SRR will be a small "perch-and-stare" system. In addition, the Army is developing the concepts for the Soldier-borne sensor (SBS) system, which is targeting a <250-g platform as a situational awareness tool, providing the squad with an organic "quick look" capability for near-real-time video feeds of larger, complex, and restrictive environments. The goal of the SBS is to reduce exposure to potential threats and enhances freedom of maneuver by providing actionable information to support decisions at the lowest echelon. Many of the capabilities and new understanding developed under the MAST CTA will directly impact future Army systems such as the SRR and SBS.

1.7 References

1. Unmanned systems integrated roadmap, FY2013–2038 version. Reference Number: 14-S-0553.
2. Unmanned ground systems roadmap. July 2011. Fort Leonard Wood (MO): Robotics Systems Joint Project Office.
3. Army robotic and autonomy systems strategy. March 2017. Fort Eustic (VA): Maneuver, Aviation, and Soldier Division Army Capabilities Integration Center US Army Training and Doctrine Command.

2. Mobility, Control, and Energetics Thrust

Thrust Lead: J Sean Humbert, University of Colorado
ARL Lead: Chris Kroninger, Vehicle Technology Directorate

2.1 Vision of the Mobility, Control, and Energetics Thrust

The MAST consortium worked toward creating systems of diverse autonomous mobility platforms equipped with miniature sensors to quickly, quietly, and reliably explore and find targets in the MAST scenarios. These platforms can exploit size (tradeoffs between payload, minimum exploitable opening, speed, range, and duration), diversity of locomotion (flight and ambulation), multiple units (redundancy, communication network, observation angles), and diversity of sensing (passive and active, local and long range). The vision of the MAST consortium was that a diverse combination of handheld mobile platforms can perform the MAST scenario tasks with low cost and high probability of success. The MCE thrust supported the mobility aims of the MAST consortium by striving to achieve this vision:

> Develop the fundamental science, tools, and algorithms that will enable the mobility of heterogeneous teams of autonomous micro-platforms for tactical situational awareness in MAST scenarios.

Effectively supporting tactically relevant dismounted operations requires that MAST-scale platforms be able to negotiate Army relevant environments (i.e., complex and dynamic). Research activities of the MCE thrust sought to provide understandings and tools that translate into robotic capability such as flight in strong winds and gusts and ground locomotion on complex terrain including through sand and leaf litter, over rocks, in tight spaces, and up vertical surfaces. In many instances, these environments must be traversed at high speeds to match the tempo of Army operations. In pursuit of longer-duration missions, novel and miniature energy sources and power transmission concepts were considered as were technologies to support efficient use of energy including behaviors like perching as well as mechanics simulation and manufacturing technologies that result in well-integrated efficient platforms.

2.2 Technical Challenges

An enduring question throughout the program was how principles of biological locomotion can be discerned and employed as inspiration for robotic mobility. Challenges include eliciting animal behaviors in controlled setting, and collecting dynamic data and distilling understanding to basic principles for integration into

exemplar platforms for behavior validation. Model development accounting for vehicle-medium (air and ground) interaction also presented challenges deriving from the nonlinear, unsteady, and stochastic nature of the medium dynamics. Fluid modeling requires simulating nonlinear, unsteady low Reynolds number flows but effective flight control demands the development of reduced order or otherwise real-time models. Relevant terra-mechanics models of the varying and often heterogeneous composition of natural surfaces were required to be developed essentially from scratch. Material processing and platform integration challenges were associated with the creation of novel high energy/power density fuel sources as well as creating small yet sophisticated robotic platforms. Fast body dynamics and high operational speeds on often intrinsically unstable platforms required closed-loop feedback at high bandwidth demanding rapid data acquisition and estimation and fast learning to respond to evolving environments. The program's explicit imposition of constraints on SWaP challenged developments across the thrust.

2.3 Background

At the inception of the program in 2008, handheld or even man-portable robotic platforms were a novelty at best and low functioning. Deriving from a history of model aircraft building, fixed and rotary wing aircraft did exist as remote controlled platforms including, for example, the AeroVironment (AV) Black Widow (Fig. 2-1). Some novel rotary designs and understanding of small aircraft flight had come from the recent Army Multidisciplinary University Research Initiative (MURI) entitled "Micro Hovering Air Vehicles: Revolutionary Concepts and Navigational Advancements". The tremendous popularity of the now commonplace quadrotor was yet to come. Principles of biological flight were being actively studied (e.g., that the leading-edge vortex was an important feature of unsteady insect flight was known but the physical mechanisms behind its persistence were still hotly debated). Some examples of flapping-wing cruising vehicles existed and the extremely novel Harvard microfly had just made its first tethered flight but it would not be until 2011 that the AV's Hummingbird would become the first tailless hovering flapping aircraft to perform untethered flight.

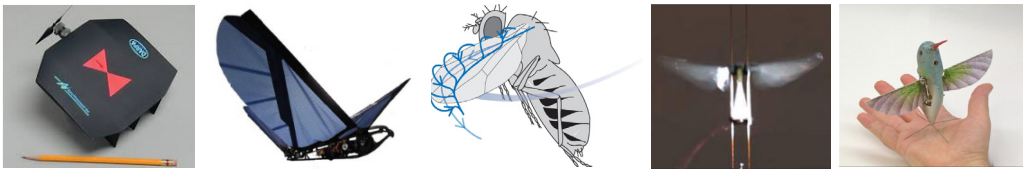


Fig. 2-1 Small aircraft and understanding developed by the inception of the program including (left to right) the AV Black Widow, Morpheus Lab ornithopter, depiction of a leading edge vortex, and the Harvard microfly. The later AV Hummingbird's flight came a bit later (far right).

The state of the art in ground mobility was represented by wheeled and tracked vehicles. Soldier-packable platforms at the time included the iRobot PackBot and small unmanned ground vehicle (SUGV). Of similar scale was the legged research platform, RHex. While smaller platforms with "whegs" provided researchers a simple wheel-leg hybrid for negotiating complex surfaces (Fig. 2-2), a more bio-inspired hexapod robots from University of California (UC) Berkeley featured an alternating tripod gait. Basic principles of legged mobility were understood at the time, including a seemingly fundamental scaling law for leg stiffness with body size. Understanding of how to adhere to smooth surfaces was gleaned from the study of geckos and the RiSE robot, which would ultimately scale smooth surfaces, was climbing rough surfaced buildings using spines. In general, small legged robots of the time were slow and executed continuous open-loop walking/running/ climbing with minimal ability to traverse complex terrain.

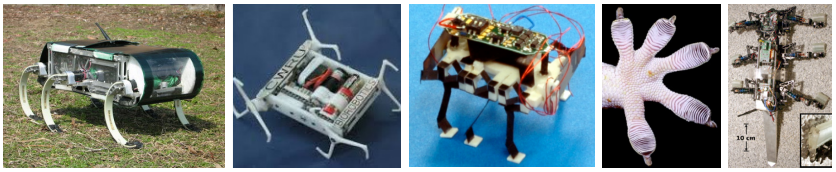


Fig. 2-2 Small legged systems of the time including (left to right) RHEX, Case Western Reserve University's Wheg platform, the RoACH platform, a gecko foot adhering to glass, and the RiSE robot ascending a stucco surface

While the MCE thrust could not bring sufficient resources to bear to meaningfully contribute to traditional battery development, the thrust was able to leverage improvements in lithium-based battery developments driven by the mobile electronics industry. At the inception of the program, research investments in novel sources of high power and energy density sources, which would be alternatives to batteries for small MAST systems, were limited with one example being the Massachusetts Institute of Technology (MIT) micro-turbine. For lower power demands like sensing, scavenging technologies such as conformal solar cells and mechanical vibration energy harvesting were being studied and developed.

A significant gap existed between the ability to build structures and integrate actuators at the macroscale, where nut and bolt assembly is the paradigm and actuators tend to be rotary, and the microscale with 2.5-D layer processing with bending or hinged actuators. In the mesoscale regime spanning this gap, fabrication approaches that allowed for the assembly of the Harvard microfly were in their infancy. It was understood that the primary driver of the large amplitude flapping kinematics in insects was two pairs or orthogonally oriented muscles changing the shape of the thorax and dielectric elastomer actuators (DEAs) were considered candidates for developing novel muscle-like linear actuators (Fig. 2-3).

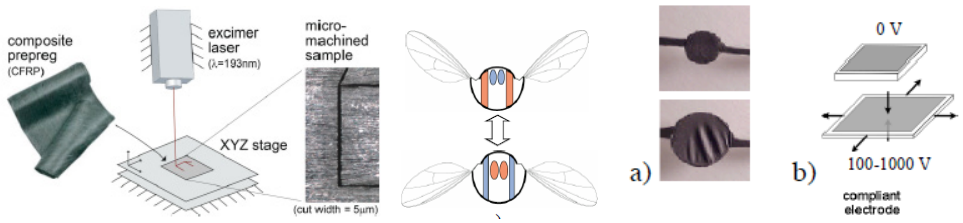


Fig. 2-3 Manufacturing and actuator technologies including (left to right) mesoscale machining and assembly approach, an illustration of muscle action in a thorax, and an a) image and b) schematic of the function of a DEA

Advancements in miniaturized electronic component including microprocessors, cameras and ranging vision, inertial measurement units (IMUs), radios, and so on, along with the availability of well-integrated microprocessors, autopilots, and cell phone chipsets enabled the realization of such small platforms from the beginning of the program. These enabled the implementation of algorithms of modern control and autonomy to be applied to MAST systems.

2.4 Thrust Organization

By the midpoint of the program, the MCE thrust was organized around four technology subthrusts: 1) Aerial Mobility, 2) Ground Mobility, 3) Actuation and Power, and 4) Dynamics and Control (Fig. 2-4). The thrust model was an organizational construct added to the program to foster collaboration across the individually proposed MAST centers (Mechanics, Micro-electronics, Processing for Autonomy, and Integration) and all the centers played a role in the MCE thrust.



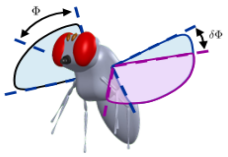
Fig. 2-4 Exemplars of the 4 MCE sub thrusts including (left to right) a cyclocopter (Aerial Mobility), legged robot (Ground Mobility), the conceptual design of a novel energy source (Actuation and Power), and a dynamic quadrotor (Dynamics and Control).

2.5 Technical Program

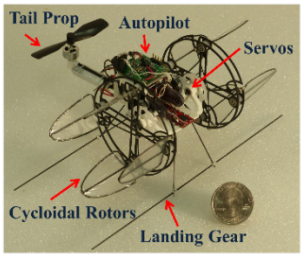
The following are summaries of the research conducted under each subthrust.

2.5.1 Aerial Mobility

Principles of Flight Derived From Biology: Taking advantage of the model simplification afforded by wing stroke averaging, linear time invariant (LTI) models of flapping-wing flight were developed from data generated in high-fidelity aerodynamic simulations (IBINS) and dynamically scaled robotic experiments (robofly). The kinematics used for the hovering and maneuvering cases were measured using high-speed photogrammetry and Kalman filtering techniques. Areas that controlled the insects visual field and disturbance environment and where animal kinematics to be measured at high speed allowed hypotheses about species flight control strategies to be explored, elucidating potential control approaches to enable micro air vehicles (MAVs) to take advantage of the full flight envelope.



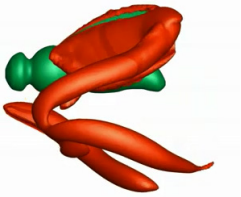
Insects have a high number of control degrees of freedom and their effects are not oriented along the orthogonal axes of a traditional vehicle state description.



A 29-g mesoscale cyclocopter.

Novel Configurations: Novel flying platforms such as cyclocopters and flapping wings were developed. Pioneering experimental (force/flow field measurements) and computational (CFD/aeroelastic analysis) studies conducted on cycloidal rotors in both hover and forward flight not only improved the understanding of this novel concept, but also helped in formulating a set of design principles for an efficient cycloidal rotor. Along with innovative vehicle design, novel flight control strategies and micro-electronics led to the development of the first flying cyclocopter. The experimental results also showed that a cycloidal rotor could be efficient, capable of flying at high advance ratios (using thrust vectoring) and gust-tolerant. Additionally flapping wing platforms at 2 different scales, designed with hummingbird-inspired attitude control mechanisms, were built and flight tested.

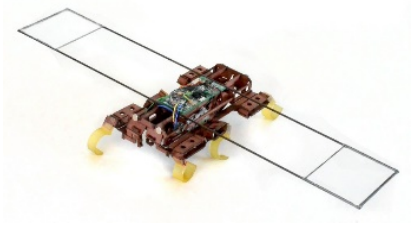
Aerodynamic Understanding and Prediction: Computational and experimental techniques have been developed and applied to the uniquely low Reynolds number unsteady flows prevalent for handheld systems. High-fidelity CFD codes were developed to perform direct numerical simulation on stationary grids and CFD codes were coupled with highly flexible membrane structural dynamics codes. CFD Adjoint optimization codes were developed for unsteady flows for kinematic and wing shape design and reduced order decompositions of high-fidelity computational and experimental data resulted in control tractable prediction speeds. This numerical work was validated with particle image velocity (PIV) of flow fields and digital image correlation (DIC) of structural deformation measurements as well as force measurement.



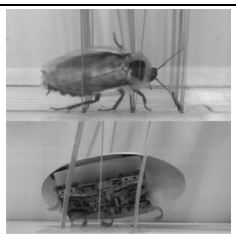
Reynolds averaged Navier Stokes (RANS) simulation of a hovering in fruit fly.

2.5.2 Ground Mobility

Steady Locomotion in Benign Environment: Bio-inspired approaches guided the demonstration of VelociRoACH. VelociRoACH runs with a spring-loaded inverted pendulum (SLIP) gait. Choosing the relative stiffness of legs to match small animals such as cockroaches, this 10-cm-long, 30-g hexapedal millirobot is capable of running at 2.7 m/s. Complementary analysis using dynamic systems theory extended our understand of the stability characteristics of hybrid gaits, which have both grounded and aerial phases.



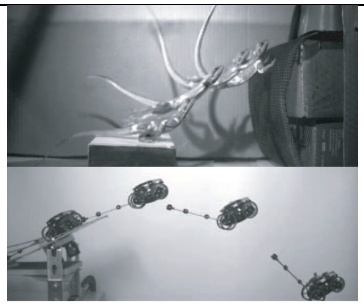
Bio-inspired VelociRoACH hexapod designed with dynamic similarity to cockroach. The vanes stabilize the roll axis.



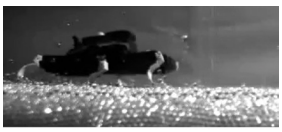
Snapshot of a cockroach and a robot with an ellipsoidal shell rolling their body to the side and maneuvering through gaps.

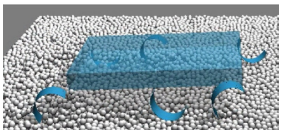
Extended Terrain Mobility – Beyond All Terrain: Cockroaches possessing thin, rounded bodies were challenged to traverse tall, narrowly spaced, vertical, grass-like structures. Animals used a novel roll maneuver, a form of natural parkour, allowing them to rapidly traverse obstacle gaps narrower than half their body width. Inspired by this discovery, a rounded exoskeletal shells was added to a legged robot. Without adding sensory feedback or changing the open-loop control, the rounded shell enabled the robot to traverse beam obstacles with gaps narrower than the shell width via body roll. Terradynamic "streamlined" is just one of many principles discovered that allow small-brained creatures to successfully navigate complex environments.

Transitions – Beyond Single Behaviors: Application of a control-theoretic framework advanced understanding of inertial appendage use in locomotion. It was discovered that lizards control the swing of their tails in a measured manner to redirect angular momentum from their bodies to their tails, stabilizing body attitude in the sagittal plane. Inertial control of tail and body has been explored for self-righting and impulsive turning in animals. Theory has been developed to describe zero-net angular momentum maneuvers, which have been demonstrated in small robots.



By swinging tails upwards, lizards redirect angular momentum from the body to the tail to maintain body angle. A small robot doing the same.

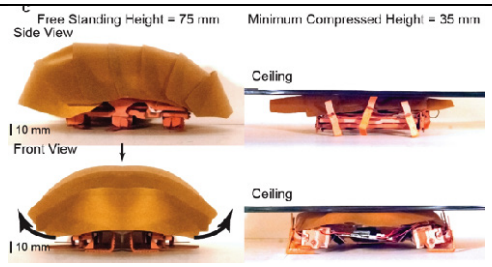




Robot moving across glass beads (top) and simulated robot in MBDyn (bottom).

Modification – Transforming Environment: The locomotion of small animals and robots moving on and within dry granular media has been systematically studied experimental. Fluidized bed trackways were developed to enable control of initial conditions of the granular medium, including compaction and inclination. The fluidized bed trackway was equipped with motor/vision systems and a universal jamming gripper, enabling systematic study of robot locomotion on heterogeneous granular media. Limb and gait design was advanced by observing animals and robots moving through flowing media. High-fidelity coupled multibody (MBDyn) and multiparticle computer simulations (using the Discrete Element Method) to model interaction forces at the particle and stress network level were developed. Resistive Force Theory (RFT), a much lower-fidelity model of ground-limb interaction, was also shown to accurately model the locomotion performance and forces generated during limb and body intrusions, including for legged locomotion, hopping, and sand-swimming.

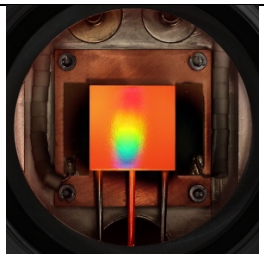
Robustness – Integrated Persistence: Jointed exoskeletons permit rapid appendage-driven locomotion while retaining the soft-bodied, shape-changing ability to explore confined environments. Cockroaches were challenged to pass through horizontal crevices smaller than a quarter of their standing body height. Cockroaches rapidly traversed crevices by compressing their body 40%–60%. To define the limits of body compression in confined spaces, dynamic compressive cycle tests were performed on living animals. Cockroach exoskeletons provided biological inspiration for the manufacture of an origami-style, soft, legged robot that can locomote rapidly in both open and confined spaces.

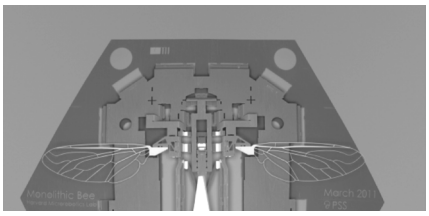


Prototype robot with adjustable sprawl and abdominal compression-inspired exoskeletal plate-like shell. Top row photos: Side view of freestanding and confined-space posture for robot between two surfaces with ceiling labeled. Bottom row photos: Front view of freestanding and confined-space posture for robot between 2 surfaces.

2.5.3 Actuation and Power

Novel Power Sources: Technology has been developed that efficiently harnesses the energy content of hydrocarbon fuels in a volume only a fraction of a cubic inch. A propane-fueled microcombustor heats a photonic crystal emitter to incandescence and the resulting spectrally confined thermal radiation drives low-bandgap photovoltaic (PV) cells to generate electricity. This approach, based on thermophotovoltaics (TPV), enables a higher energy density storage than batteries because hydrocarbon fuels have 60 times the energy density of lithium-ion batteries. To support sensor payloads for persistent stare operation, thin membrane solar cells have been developed. These can be integrated into structures like wings.

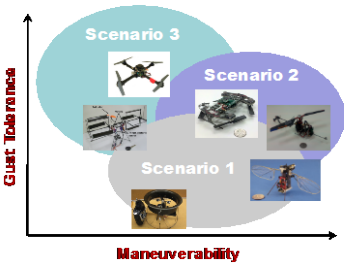

Microcombustor with integrated photonic crystal (diffraction pattern visible) in operation.


Pop-up book assembly of the Harvard microfly.

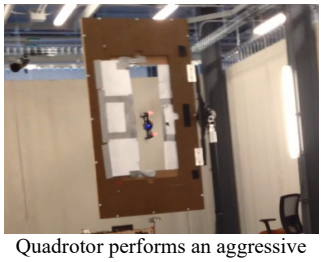
Micro-Manufacturing: Smart composite-based rapid manufacturing techniques were developed that control assembly and functional degrees of freedom during fabrication, achieving the automated assembly of complex devises using “pop-up book” techniques. This approach allows for mechanical alignment of bulk-machined MEMS devices, a single assembly degree of freedom, and the capability to assembly monolithic devices (including actuators and electronic components).

2.5.4 Dynamics and Control

System ID, Reachability, and Bio-Inspired Perception: Tools from dynamical systems and system identification theory were applied to high-fidelity simulations and open-loop flight data for a variety of small-scale platforms to generate flight dynamics models. In particular, norms of reachability and disturbance sensitivity gramians were used to characterize the maneuverability and gust rejection capability of candidate small-scale platforms, allowing for a mapping of the platforms into MAST scenarios.



Quantification of the bare airframe maneuverability and gust tolerance for various aircraft.



Quadrotor performs an aggressive maneuver with onboard vision and trajectory planning.



Quadrotor with an array of accelerometers.

Behaviors and Perception: Two perspectives have been explored for enabling small vehicle to perform dynamic behaviors in complex and high disturbance environments, one guided by computer science and algorithm development and the other bio-inspired. For example, abstractions of quadrotor dynamics allowed for the development of low degree of freedom to nonlinear flight controllers, enabling path planning for aggressive trajectories to be derived from Newtonian mechanics allowing flight through confined spaces or to perches and grasp objects. Additionally, two bio-inspired sensorimotor feedback approaches were investigated and validated in flight experiments. Signals from arrays of distributed accelerometers were processed to provide the three translational accelerations and angular acceleration components, which were subsequently mapped through mass and inertia to provide estimates of external force and moment on the vehicle. When applied as feedback, these signals provide a significantly improved gust/disturbance rejection capability. As regards vision, the biological principles for two proposed pathways in the insect visuomotor system for detecting small objects were characterized and implemented on flight vehicles, demonstrating local perception in cluttered environments with very low SWaP requirements.

Understanding and Learning: Mathematical and computational tools were developed to learn the underlying uncertainty that affects the behavior of MAST-scale platforms. The developed tools offer probabilistic guarantees of validity, and were tested in legged and wheeled robots operating under uncertainty and aerial robots operating under the influence of uncertain aerodynamic (ground) effects. Comprehensive experimental study contributed to better understand the impact, model, and adapt to ground effects in high-speed flight.

The field of machine learning has had significant impact on robotics, particularly on perception. More recently, we have been able to adopt deep learning architectures with memory—the long short term memory (LSTM) and differential neural computer (DNC)—to learning for planning and control, without a priori constructing a global map. Extensive training in simulated environments has yielded principles of transfer learning from autonomous agents in simulation to actual robotic vehicles.

2.6 Significant Accomplishments

Over the course of the past decade, the MCE thrust has advanced the mobility capabilities of small-scale platforms to support operation in Army-relevant environments. Since a complete list of accomplishments would be onerous, we have reported here a few highlights of key technical accomplishments:

1) *Advancements of novel platforms*: Novel platform developments include demonstrating the first successful flight test of a cycloidal rotor-based aircraft. Other novel platforms include the high-performance jumping robot Salto, several variants of the Hexa/OctoRoACH including the high-speed VelociRoACH and the winged bipedal ornithopter (BOLT).

2) *Perception and motion control*: Visual-inertial fusion has moved flight control from off-board computers coupled to motion capture systems to onboard intelligence. Nonlinear controller design has allowed quadrotors to plan and execute rapid aggressive maneuvers while bio-inspired principles have enhanced the perception of gusts and small obstacles supporting robust maneuvering. A developed gramian-based framework provides a formal tool to design control system degrees of freedom to achieve desired control responses.

3) *Discernment and instantiation of aerial behaviors*: We addressed experimental and simulation trade studies to quantify controller changes to improve estimation and adaption to gusts as well as trajectory generation for improved gust rejection. We addressed fast, online generation of motion plans for MAVs that are informed by capabilities of UAVs subject to gusts. We have also addressed dynamically perching on and grasping of a variety of surfaces and reasoning for understanding about failed perching attempts.

4) *Discernment and instantiation of ground-based behaviors*: A particularly important advancement is the appreciation of the function of inertial appendages (e.g., tails) for transient control. Tails have enable robots to turn rapidly and self-right. Other important behaviors include walking, running, and crawling across various surfaces, including flowing media, and through confined spaces.

5) *Experimental design for transient and discrete interactions*: Ensuring resilience required in real-world environments requires negotiating unsteady interactions. A variety of arenas have been developed that can introduce well-characterized wind gust repeatably in well-controlled environments. Discrete interactions of legs in complex environments can now be explored on an apparatus built on a fluidized bed with computer-

controlled obstacle deployment. These and other test beds have served the collection of important data sets of animal and robotic locomotion.

6) *Mesoscale manufacturing techniques*: The smart composite manufacturing approach matured during the program. This technique allowed for the automated repeatable construction of small complex mechanical systems like the microfly and has been used for robot prototyping throughout the program.

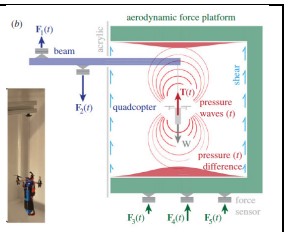
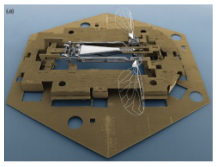
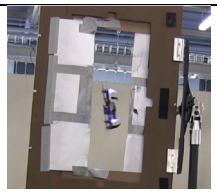
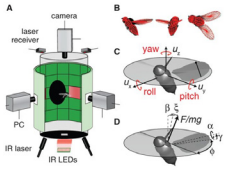
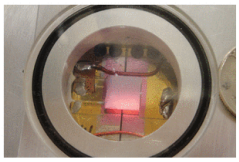
7) *Models for motion*: Many fluid and terradynamic models have been created for a variety of purposes including detailed models for understanding fundamental physics to lower fidelity approximations to deploy for real-time motion control. A direct numerical simulation (DNS) CFD code was developed and compared to existing higher Reynolds number approaches, while optimization and model reduction was explored with adjoint models and proper orthogonal decomposition. Terramechanics models included the development of resistive force theory, a very low fidelity model that produces results surprisingly similar to higher fidelity discrete element modeling (DEM). Both fluid and terradynamic models were coupled to novel nonlinear structural models for extreme deformation.

8) *Thermo-photovoltaic electricity source*. This novel power source demonstrates record heat-to-electricity conversion at the mesoscale and represents a viable future alternative to energy limited lithium-based batteries for high power density needs.

9) *Power-aware operation of teams*: While interested in op-tempo deployment of teams of quadrotors in urban environments, we advanced the energetics of deployment by benchmarking of component-level performance in quadrotor aircrafts and performance across different length scales.

2.7 Metrics of Success

2.7.1 Some High-Impact Publications

Salto robot achieved continuous vertical jumping agility of 1.7 m/s. Cover of 1st *Science Robotics* issue.[1]	Demonstrated the change in energy storage and output for a lizard running on hard and flowing media. Cover of Experimental Biology[2]	Inertial assisted control – new field in biology and bioinspired robots providing simple control of transitions through appendage inertia. Cover of Nature [3]	 Developed non-intrusive aero-dynamic force measurement in freely flying animals and robots.[4]
 Pop-up box fabrication of a fly-scale flapper. [5]	 A 250-g quadrotor performs aggressive maneuvers without a motion capture system. [6]	 Explores predator evasion strategies in flies. [7]	 Incandescent emitter powered by catalyic combustion. [8]

2.8 Challenges for the Future

The primary challenges for MAST-scale robotic mobility continue to derive from the need to move through complex environments (e.g., subterranean, mountainous, megacity, wooded/arboreal), that is, terrain that impedes Soldier mobility. The limited access to high energy density fuels will also limit performance. Unlocking more advanced locomotive capabilities awaits advancements in low SWaP, perception, and intelligence. While physically mediated feed forward control and simple feedback strategies have taken small robot surprisingly far, moving more rapidly and efficiently through complex spaces will required greater cognitive function; perception and intelligence allow small robots to path plan, learn, and practice.

To be positioned to exploit higher cognitive and control functions, advanced materials, material processing, and robotic integration is required to developed more sophisticated platforms with higher degrees of freedom. Advances in high

power density actuators of various form factors integrated into local physical networks with sensing, computation, and energy storage will expand the range of motion, allowing small robots to be more adaptable, whether this design sophistication is used to create distributed reflexive responses to disturbance or create gross configuration changes to support different modes of operation.

Access to energy will also continue to limit MAST-scale robotic performance. Batteries, with much lower energy densities than larger-scale hydrocarbon fuel systems, remain the best choice at these scales. The handheld electronics market will continue to be a driver for high power density rechargeable batteries and small robots can leverage these advancements, but rapid or disruptive advances are not likely. Alternative sources for energy may have a role in the future of robotic mobility. Using energy available in the environment will require the exploitation of energy sources that offer a substantial fraction of the mobility budget. This generally rules out solar energy for MAST-scale systems, but wind and biomass might prove viable sources.

2.9 MCE Task Summary

CENTER LEAD: Sean Humber, University of Colorado, Inder Chopra, University of Maryland			
ARL TECH LEAD: Chris Kroninger			
Project #	Project Title	PI	Org
MCE-16-1	Aeromechanics Subthrust	Humbert	Maryland
MCE-16-1.1	GPU-Based CFD Analysis For Rotary- and Flapping-Wing- Based MAVs	Baeder	Maryland
MCE-16-1.2	Bio-Inspired Tandem-Wings: Kinematics, Flow Physics, Gust Tolerance and Maneuverability for MAST-Scale Aerial Platforms	Laurence	Maryland
MCE-16-1.3	Highly-Maneuverable, High-Speed, Optimized Next-Generation Micro Cyclocopter	Chopra / Benedict	Maryland, Texas A&M
MCE-16-1.4	Robust Comprehensive CFD/CSD Methodology for Multidisciplinary Design Optimization of MAST-Scale Aerial Vehicles	Chopra / Masarati	Maryland
MCE-16-1.5	Reduced-order Model and Shaped Sensor for Flapping-Wing Control stability of a flapping wing	Wei / Sirohi	New Mexico
MCE-16-2	Ambulation Sub-Thrust	Full/Fearing	Berkeley
MCE-16-2.1	Robophysical approach to discovery of principles of effective interaction of ambulatory robots in multicomponent environments	Goldman	GaTech
MCE-16-2.2	Bio-inspired Principles of Extended Terrain Mobility	Full	Berkeley
MCE-16-2.3	Bio-inspired Principles of Transitions	Full	Berkeley
MCE-16-2.4	Real-time Force Feedback in Heterogeneous Terrain	Fearing	Berkeley
MCE-16-2.5	Dynamic Locomotion in Urban Environment	Fearing / Full	Berkeley
MCE 16-3	Actuation and Power	Celanovic	MIT
MCE-16-3.1	Ultra-high energy density TPV generator for small robotic platforms : First ever demonstration of fuel powered robot with extreme range	Celanovic / Joannopoulos	MIT
MCE-16-4	Dynamics and Control Sub-Thrust	Humbert	Colorado U
MCE-16-4-1	Bio-Inspired Sensorimotor Control for MAST-Scale Aerial Platforms	Humbert	CU-Boulder
MCE-16-4.2	System Identification of MAST Scale Platforms	MacFarlane / Chopra	Maryland
MCE-16-4.3	Unraveling Avian Gusts Mitigation Strategies to make MAVs more Robust	Lentink	Stanford
MCE-16-4.4	Integrated Air-Surface Operation for Micro Air Vehicles	Cutkosky / Humbert	Stanford / CU-Boulder
MCE-16-4-5	Landing, perching, and grasping for micro aerial vehicles	Kumar / Daniilidis / Sreenath	Pennsylvania
MCE-16-4-6	Controlled Mobility for Multi-Robot Risk-Tolerant Reconnaissance	Tanner	Mexico
MCE-16-4-7	Learning to Adapt Mobility to 3-D Unstructured Environments	Abbeel	Berkeley

2.10 References

1. Haldane, D. W., Plecnik, M. M., Yim, J. K., and Fearing, R. S. (2016). Robotic vertical jumping agility via series-elastic power modulation. Science Robotics, 1(1).

2. Chen Li, S. Tonia Hsieh, and Daniel I. Goldman. (2012). Multi-functional foot use during running of the zebra-tailed lizard (Callisaurus draconoides). J. Experimental Biology, 215, 3293.

3. Libby, T., Moore, T. Y., Chang-Siu, E., Li, D., Cohen, D. J., Jusufi, A., and Full, R. J. (2012). Tail-assisted pitch control in lizards, robots and dinosaurs. Nature, 481(7380), 181-184.

4. Lentink D, Haselsteiner AF, Ingersoll R. (2015). In vivo recording of aerodynamic force with an aerodynamic force platform: from drones to birds. J. R. Soc. Interface.

5. Sreetharan, P. S., Whitney, J. P., Strauss, M. D., and Wood, R. J. (2012). Monolithic manufacture of millimeter-scale machines. J. Microengineering Micromechanics, 22(055027).

6. G. Loianno, G., Brunner, C., McGrath, G., and Kumar, V. (2017). Estimation, control and planning for aggressive flight with a small quadrotor with a single camera and IMU. Int'l Conf on Robotics and Automation.

7. Muijres, F. T., Elzinga, M. J., Melis, J. M., and Dickinson, M. H. (2014). Flies evade looming targets by executing rapid visually directed banked turns. Science, 344(6180), 172-177.

8. Chan, W. R., Bermel, P., Senkevich, J., Ghebrebrhan, M., Yeng, Y. X., Marton, C. H., Jensen, K. F., Johnson, S. G., Joannopoulos, J. D., Soljacic, M., and Celanovic, I. (2013). High energy density, high-efficiency, moderate-temperature chip-scale thermophotovoltaics. Proceedings of the National Academy of Sciences, 110(14), 5309-5314

GPU-Based CFD Analysis for Rotary- and Flapping-Wing- Based MAVs

Project: MCE-17-1.1

James D. Baeder, Associate professor
Department of Aerospace Engineering
University of Maryland
College Park, MD
Baeder@umd.edu

Camli Badrya, Postdoctoral associate
Department of Aerospace Engineering
University of Maryland
College Park, MD
cbadrya@umd.edu

Christopher M. Kroninger,
Aerospace engineer
Army Research Laboratory
Aberdeen Proving Ground, MD
christopher.m.kroninger.civ@umd.edu

***Abstract*— From an engineering perspective, it is useful to compare the performance of different flapping wing configurations (in steady forward flight or hover) under ``trimmed'' flight conditions with themselves as well as other micro-air-vehicles. Without an efficient trim algorithm, trial-and-error based identification of the trimmed wing kinematics is computationally expensive for any flight condition, because the large number of simulations required make the process practically infeasible. In a global sense the nature of forces produced by flapping wings closely resemble those on a helicopter blade, such that an analogy can be drawn between the two. Therefore, techniques developed for helicopter performance calculation are adapted and applied to the flapping wing platform particularly for analyzing steady flight. The aim of this study is, therefore, to formulate a computationally efficient and robust framework to obtain trim solutions that couples a flight dynamic model using simplified quasi-steady (QS) wing aerodynamics to a high-fidelity CFD analysis. Such an analysis will provide the ability to incorporate the effects of detailed aerodynamic forces from CFD into the trim solution for flapping wing kinematics for the first time. This work also highlights the need to develop efficient computational facilities and take advantage of GPU technology.**

(key words: Flapping-wing, trim, viscous flow, MAV)

I. Introduction

Background and Technical Challenges:

CFD-based multidisciplinary aeroelastic MAV simulations may: (1) function as an accurate performance prediction tool, (2) improve understanding of fluid flow physics and vehicle stability, (3) act as test bed for simulating new design concepts, and (4) support development of efficient reduced order models that can be utilized in routine design. The main physical and computational challenges for MAV simulations are related to: the relatively low Reynolds numbers for which small changes in the Reynolds number may result in significant changes in the flow physics; gridding issues related to capturing accurately the boundary layer flow in the vicinity of walls as well as the off surface vertical structures and acoustic waves; the necessity of CFD/CSD coupling for flexible structures; the large range of spatial and temporal scales present in the flow; and the computational expense, especially for cases where the vehicle may be undergoing maneuvers or encountering unsteady gusts.

The aerodynamic mechanisms of flight at low Reynolds numbers and their impact on stability present significant physical modeling challenges. In the case of insect flight, massive separation causes the creation of a leading edge vortex structure and may interact with the clap and fling mechanism. Initially laminar flow can transition to turbulence as a result of complicated vortex dynamics. Thus, 3-D unsteady Navier-Stokes equations are required. Reynolds numbers for MAVs vary from 10^2 for insect-based flight to 10^4-10^5 for avian-based and rotor-based flight. Direct Numerical Simulations (DNS) can be used in the low Reynolds number regime, where the flow is mostly laminar and transitional. For cases in which the flow is turbulent, fully resolved DNS is possible for Reynolds numbers up to 10^3. Large Eddy Simulations (LES) can provide significant savings in mesh requirements (over DNS) in wake and near-wall regions. Thus, LES should be possible for Reynolds numbers approaching 50000. For calculations at the higher end of the spectrum, a hybrid RANS-LES (e.g. Improved Delayed Detached Eddy Simulations - IDDES) approach can be utilized, with RANS near the solid surfaces and LES in the wake regions. For engineering calculations at the higher range of Reynolds number the Unsteady Reynolds Averaged Navier-Stokes (URANS) can be used.

Some of the remaining technical challenges that remain in applying CFD to MAV simulations are to: improve the overall efficiency of the CFD simulations; extract from CFD the aerodynamic and structural power required for the various MAV in flight; obtain trimmed flight in hover and forward flight (by changing the flapping kinematics); and extract reduced order models from the detailed CFD simulations to aid in optimization. At the same time CFD needs to be used to examine the most efficient configuration to provide required MAV capabilities.

Research Objects:

To increase the sophistication of physics-based CFD simulation tools of highly maneuverable flying platforms for: (1) understanding the complex aerodynamic environment of current and future MAV systems, (2) validating with experimental data, (3) looking at static stability and control derivatives, (4) trimming through the change in flapping kinematics, (5) handling unsteady maneuvers and gusts, and (6) coupling to comprehensive aeromechanic analysis. This

should be performed in a manner that is consistent with spreading this technology as rapidly as possible to other tasks within the CTA MAST. Thus, a focus during the three year period is to rapidly assemble and take advantage of GPU-based CFD simulation capabilities, both at the workstation level and at the HPC level.

II. TECHNICAL APPROACH

In order to carry out the four scenarios of progressive difficulty (small unit building search, small unit cave or collapsed building search, small unit perimeter defense, and small unit mobile real time ISR) there is a requirement for physics-based simulation tools to enable first principles-based understanding of the aeromechanics at the small scales to enable the design of agile, efficient, gust tolerant aerial platforms. The primary innovation in the CFD simulation methodology is to take advantage of GPU-based algorithms to provide a one to two orders of magnitude increase in the level of compute capability over traditional CPU-based techniques. This enables the simulation of unsteady gust interactions and more complex configurations of interest to MAST systems.

Curvilinear overset methods are more suitable and proven for rotor-based (e.g. cyclocopters [1] and shrouded rotors [2]) as well as avian-based vehicles [3]. They provide the advantage of using body conforming meshes improving the accuracy of flow structures near the thinner boundary layers that occur near solid boundaries for higher Reynolds number flows [4]. Recently, multidisciplinary analysis platforms, utilizing CFD/CSD coupling, were devised, tested and validated for rotary wing vehicles [5,6] and flapping avian wings [7]. Systematic investigation though is still required to fully validate these platforms for modeling the aerodynamics and structural dynamics of flapping MAVs.

To accurately resolve multiple spatial and temporal scales associated with unsteady separated flows around MAVs and to minimize the accumulation of the discretization error in time-dependent MAV flow simulations, high-order discretization in time/space are required. New techniques for improving the spatial accuracy have recently been developed and examined at UMD to improve the numerical efficiency and spatial resolution of our CFD solvers [8,9]. Compact reconstruction (specifically CRWENO) is being examined because they have smaller stencils yet higher spectral resolution (bandwidth efficiency) than explicit reconstruction. Furthermore, they have reduced dissipation over the resolved frequencies. Finally, they can have an order of magnitude smaller error than an explicit scheme of the same order. They are particularly relevant for codes that look to resolve more of the turbulent scales (LES or IDDES). Preliminary results indicate that the compact reconstruction can be 3-5x more efficient than explicit reconstruction. Improvements to temporal reconstruction (for accuracy and efficiency) still need to be examined and implemented.

The parallelism in existing RANS solvers are still based on the domain decomposition paradigm with blocks of the mesh residing on different processors with information transferred between them at every iteration. For peak performance, these methods need to run on dozens of Intel Xeon processors. It is desired to have the capability to run these simulations efficiently on a single-desktop machine. In this regard, the use of graphics processing units (GPUs) offers the most promise. Once meant only for use in real-time, graphics rendering, these platforms have become increasingly popular in the scientific community owing to their high floating-point operation speeds and large memory bandwidth. Using NVIDIA's Compute Unified Device Architecture (CUDA), a developer is now able to write code that can execute on the GPU using a Single Instruction Multiple Data (SIMD) paradigm. For computations, several thousands of threads are spawned over hundreds of cores with each thread mapped to an array index (or an abstract entity like a pixel, finite-volume, a vortex filament, a sediment particle, etc) and capable of executing a given instruction set. This parallelism is well suited for fine-grain operations where the number of data- dependencies per computation is much lower than the size of the problem. Modest performance gains have been reported through the use of GPU technology in both structured [10] and unstructured [11-12] grid-based CFD methods. Thomas, Amiraux and Baeder [13] perform all necessary computations on the GPU platform, while using the host only for initialization and file I/O purposes. That enabled speedup of 15-50x that of a single core on the CPU, enabling fairly complicated rotorcraft brownout simulations on a single workstation.

Some of the remaining technical challenges that remain in applying CFD to MAV simulations are to: improve the overall efficiency of the CFD simulations; extract from CFD the aerodynamic and structural power required for the various MAV in flight; obtain trimmed flight in hover and forward flight (by changing the flapping kinematics); and extract reduced order models from the detailed CFD simulations to aid in optimization. At the same time CFD needs to be used to examine the most efficient configuration to provide required MAV capabilities.

Delta Coupling Approach for Trim: As stated previously, the basic suite of CFD tools have already been implemented and validated for several MAV configurations with prescribed controls/kinematics. One of the main capabilities have been developed is the capability to vary the controls/kinematics to obtain trimmed flight in hover and forward flight. During this mission period, a brute force method has been examined to obtain the first ever set of CFD simulations for trimmed MAV flight in hover for an insect with wings and body: determining a trim Jacobian (calculating the changes in forces and moments due to changes in kinematics.

GPU-based Computing: As stated previously the performing of all necessary calculations on a GPU can result in speedups of 25-50x that of a single core of the CPU, even for implicit time marching methods with implicit reconstruction (CRWENO). To enable larger simulations and additional speed up in computation times it is planned to utilize a single workstation consisting of 4 Titan GPUs. The total of 24GB of VRAM should allow for simulations on the order of 20-25 million mesh points at speeds of 100-200x that of a single core on the CPU. Thus, it is expected that a single workstation

should be equivalent to a small computer cluster. In addition, it is expected that next generation DOD HPC systems will be heterogeneous (DOE Titan, fastest in the world, already consists of 1000's of GPUs). To take advantage of multiple GPUs the existing single GPU version of OVERTURNS needs to be extended to incorporate efficient mesh splitting, overset connectivity. This can then be rapidly exploited by other tasks that require CFD simulation capabilities.

III. SNERGY WITH EXISTING WORKS AND EXTERNAL COLLABORATIONS

This work builds upon many of the CFD capabilities initially developed under the Army's MURI on hovering MAVs and applied to many rotary and flapping vehicles during the first years of the CTA-MAST. It also leverages initial GPU-based CFD analysis begun under AFOSR MURI on Rotorcraft Brownout. Finally, work on IDDES and laminar/turbulent transition leverages work initiated under the VLRCOE as applied to rotorcraft applications.

Extensive collaboration has taken place within the CTA-MAST (Tasks MCE-16-1.2, MCE-16-1.4, MCE-16-1.5, MCE-16-2.3, and MCE-16-3.2).

This research was actively collaborated with ARL-VTD (Singh, Kroninger and Pulskamp). ARL lab provided experimental measurements for this research. ARL-VTD was and will continue to support UMD with access to DOD HPC facilities.

IV. METHODOLOGY

A. GPU-Based CFD platform:

The solver that has been developed at university of Maryland is 3D Navier-stokes line solver for structured meshes, ("GARFIELD", GPU accelerated rotor flow field solver), which allows for multiple GPUs using MPI and CUDA. C++/C is used for internal CPU-GPU transfer. Memory pointers shared with Python for external overset communication 29% of wall-time is spent for data transfer (CPU-GPU overset communication) which will be reduced in future. GARFIELD on 2 GPUs is faster than OVERFLOW on 32 cores for the same grid.

Code	Cores or GPUs	Time (s)	Speedup over 32 cores
OVERFLOW	32 Core	328.95	1.0
GARFIELD	1 GPU	586.01	0.56
	2 GPU	296.28	1.11

A. CFD Solver:

CFD simulations to analyze flapping wing were carried out using 3-d OVERTURNS solver as a base line, and it will be described briefly in this section. In OVERTURNS, the baseline algorithm uses a finite volume approach to solve the unsteady compressible Navier–Stokes equations in space and time. For spatial reconstruction of the inviscid fluxes, a third-order Leers Monotone Upstream-Centered Scheme for Conservation Laws (MUSCL) scheme with Korens [1,2,3] limiter was used with an option to use higher order schemes such as the Compact-Reconstruction Weighted Essentially Non-Oscillatory (CR-WENO) schemes [8]. The viscous terms were computed using a second order central differencing scheme. Because the flapping-wing simulations are performed at low Mach numbers $O(10^{-3})$ and low Reynolds numbers $O(10^{2})$, Roe-Turkel low-Mach preconditioning was applied to improve the convergence and accuracy of the compressible NS solver. These low Reynolds numbers result in either a purely laminar flow or with eddies large enough that they can be captured such that no turbulence model was required.

Once the CFD solver is verified against experimental results, it can be used to provide insight about more complicated kinematics and trends that experiments find challenging to capture.

B. *Overset Grid System*:

A single structured mesh cannot preserve the complexities in the flow-field such as tip vortices, trailing edge vortices and far field wake features. The approach in the present work to account for moving meshes is the use of overset grids. The insect is modeled as a collection of two rigid flapping wings and a rigid body, the meshes of which are overset onto a coarser nested background mesh. Nested Cartesian background meshes are used to capture the far-field effects. For the background grids a simple Cartesian mesh generator is used for this purpose. Two levels of nested Cartesian background are used. The wing grids translate and rotate according to user-prescribed kinematics while the body mesh is fixed for hover and forward flight modes. Such a setup minimizes the total mesh points required to obtain a detailed resolution of the flowfield, and allows for selectively refining the flow domains of interest. The overall overset grid system has five grids for the diptera configuration, as shown in Fig. 1, and consists of roughly 10 million mesh points (depending on the simulation type). The important grid parameters are the following: a) spacing near the wall (wing surface) is 0.01% of the mean chord, and the wings grid is extended in the normal direction over 82 planes to resolve the near-wake features b) Wing thickness is 5.8% of the averaged chord.

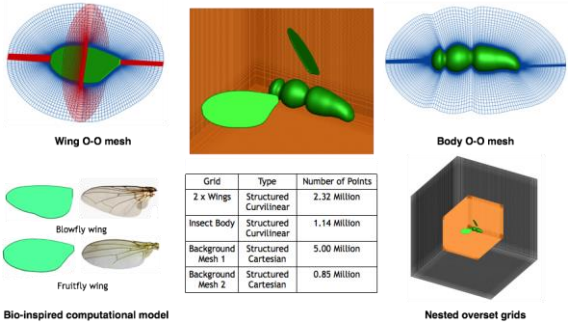


Figure 1. 2Nested overset grid system highlighting the near body wing and insect body meshes.

Body-fitted, O-O curvilinear grids are used for the wings and the insect body, which essentially capture the near-field flow, i.e., around the wings and body surface. A 3D algebraic mesh generator [18] for structured, curvilinear, body-fitted

CFD grids was used for the wings and insect body grids. The insect body and wing meshes are finely spaced near the surface to resolve the boundary layer features and associated viscous effects, as illustrated in Fig. 1.

C. *Wing Kinematics*

A simplified kinematic representation is used in the present study, which consists of three positional angles as shown in (Fig. 2) 1. Flap angle f : the wing position in the flap stroke, 2. Stroke plane angle b, the inclination of the stroke plane with respect to the x-axis, and 3. Pitch angle a_g , the angle between the wing chordwise direction and the flap-plane, also referred to as "geometrical angle of attack."

Wing kinematics used in this study are qualitatively similar to those used by Faruque and Humbert [16] for Dipteran insect flight dynamics, and by MacFarlane and Humbert [17] for Drosophila type platforms. The flap stroke velocity increases at constant rate and the maximum velocity occurs at mid-stroke. The azimuthal position of the wing in the flap cycle is given by $f(t) = -f_{\max}\cos(2\pi ft) + f_{off}$, where $f_{\max}$ is the peak-to-peak amplitude and f_{off} is the stroke bias. The geometric pitch angle with respect to the flapping stroke plane remains constant except near the stroke reversal during supination and pronation, and is given by $a_g(t) = 90 - a_{\max}\tanh(2.7\sin(2\pi ft))$, where $a_{\max}$ is the complimentary angle of the maximum geometric pitch at the midstroke with respect to flapping stroke plane. This angle remains constant during the translation periods and increases as the wing approaches stroke reversal. The wing aspect ratio (AR) is defined as wing-span squared (b) divided by the wing area. The wing shape is bioinspired by blowfly shape and fruitfly shape Fig. 1, and the wing assumed to be rigid. The simulation parameters are bio-inspired by drosophila flight features, i.e., flapping frequency range from 100–300 Hz, wing aspect ratio of 2.8, and Reynolds numbers of $O(10 - 10^2)$.

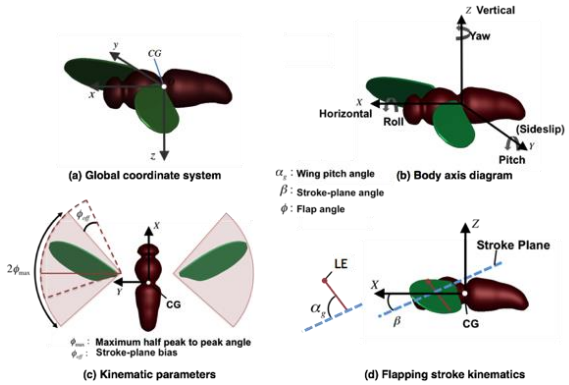


Figure 2. Inertia coordinate systems and wing kinematics parameters

V. Results and Disscussion

A. *Aerodynamic Modelling:*

Similarities in the wake can be observed between flapping and rotating wings and a comparison of the flow fields in hover is shown in Fig. 3. The rotor flow field is experimentally obtained using particle image velocimetry (PIV) and the flapping wing flow around a Diptera is obtained numerically through CFD. Note the presence of coherent tip-vortices and root vortices in the flow fields. The contraction in the wake of the flapping wing is not as pronounced as that of the rotor flow. Nonetheless, sufficient similarities exist that warrant a deeper understanding of the aerodynamics of flapping wings.

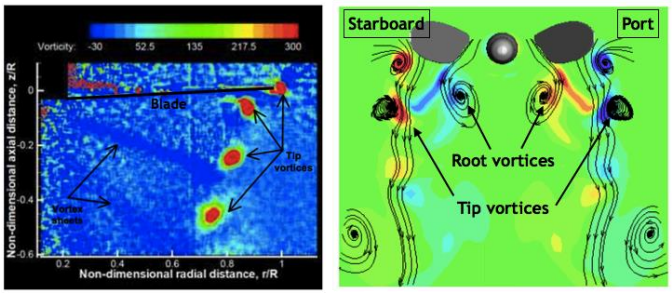


Figure 3. Comparison of flow fields in hover between a: (a) Rotor, and (b) Flapping wing. Source of experimental PIV: Ramasamy et al.

Figure 3(a) shows an instantaneous image of the vortices in the wake beneath a flapping wing in hover obtained using CFD. The signature of the tip and root vortex can be seen in the image. The vortices convect down axially and dissipate into the wake. The reason for the dissipation of the vortices are because of the natural viscosity in the flow that causes them to breakdown, and the coarsening of the meshes further away from the wing, which results in added numerical dissipation because of the increasing mesh size. Nonetheless, the signature of the flow is similar to those observed in rotor wakes in hover (see Fig. 4 (a)).

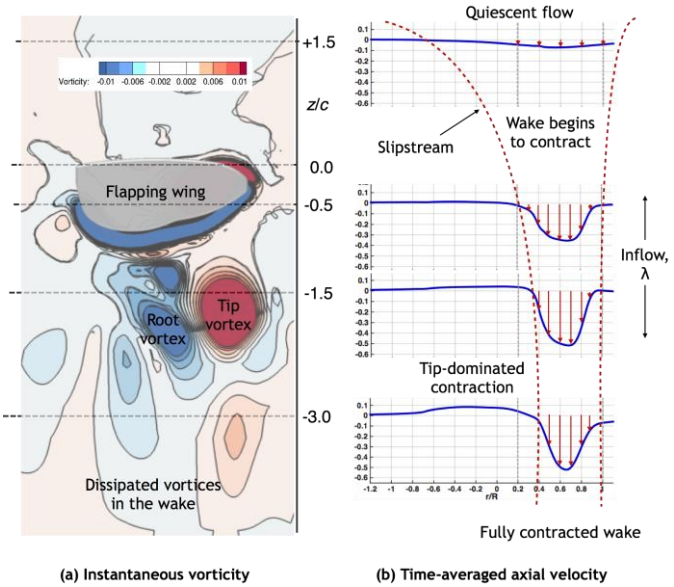


Figure 4. Schematic diagram of the hovering flap-stroke. Flap-stroke is horizontal. The shed lines describe the planes parallel to flap-stroke at different location in the Z direction below and above the wing. The location is given in nondimensional quantity (normalized Z by R).

Figure 4 (b) shows the time-average result of the axial velocity

that similarities between Fig. 4 (b) and momentum theory in hover for rotors exist. A quiescent flow upstream of the wing is observed, following a contracting of the wake through the wing, which finally asymptotes to a fully-contracted wake. The downward arrows at each longitudinal section indicate the magnitude of the axial velocity at a given spanwise station. The contracted wake is, however, more tip-dominated, which is expected for both rotors/wings that are aerodynamically loaded towards the outboard sections. While this study does not proceed to derive an appropriate contraction ratio or slipstream velocity relations, the information presented can provide certain clues to do so in future studies.

B. Trim Algorithm:

A plethora of studies exist where the kinematics are prescribed and the resulting forces and moments are studied, i.e., the direct problem. However, the interest in this section (and for practical designers) is to identify the kinematics that will result in the desired forces and moments (trim solution), i.e., the inverse problem. Solution to a trim problem requires the formulation of a Jacobian (sensitivities of the forces and moments to kinematic inputs), which is in effect a new converged CFD solution, making the process prohibitively expensive. Therefore, studies in the past have obtained trim solutions only for simplified aerodynamic models with very low computational cost. However, these simplified models do not contain the key unsteady phenomenon associated with flapping wings and as a consequence, the wing kinematic parameters are incorrect or a trial and error approach is used to only trim the lift. There is a clear need to incorporate three-dimensional CFD simulations to capture the physics of flapping wing flight. This section provides the details of the flight dynamic model, trim condition, trim variables and the loose-coupling algorithm, an idea borrowed from the world of rotorcraft where similar constraints are placed on the computational expense of computing a Jacobian. Vehicle trim is defined as the flight state when the linear and angular accelerations are zero. Within the context of flapping-wing MAVs, trim is said to be achieved if the steady state forces and moments are periodic and do not change over successive flapping cycles, i.e., the accelerations are zero in a cycle-averaged sense. This cycle variation in the forces are a result of the flap kinematics and therefore, the loads at the center of gravity of a flapping wing MAV are oscillatory in nature, fluctuating about a non-zero mean. The overall objectives of the coupled (Quasi steady/CFD) trim methodology are: 1. Generalize the flapping-wing trim framework and CFD solver to analyze steady flight including hover, forward flight and coordinated turn, 2. Determine kinematics-driven limitations for the flight envelope in level turn, and 3. Apply CFD coupling to characterize the flight limits, i.e., maximum forward speed) for a particular platform.

Flight Dynamics: The flight dynamics model is obtained by expressing the fundamental equations of force and moment equilibrium along the inertial coordinate system, shown in Fig. 2, for a general flight condition. The body of the insect is assumed to be rigid and the dynamics are given by the following equations:

$$F = \left(\frac{\P}{\P t} + W\times\right) mV + mg$$
$$M = \left(\frac{\P}{\P t} + W\times\right) \mathrm{I}\, W \quad (3)$$

where F and M represent, respectively, the external forces and moments acting at the center of gravity due to body drag, and wing aerodynamic and inertial loads. m is the total mass of the insect. At these length scales, the mass of the wings are negligible compared to body mass. g is the acceleration due to gravity and ω is the body angular velocity vector. I represents the inertia tensor with body-fixed axes as reference. Wing inertial loads are ignored for trim since they are guaranteed to have a mean value of zero for periodic flapping.

Steady flight trim represents the condition where the insect center of gravity experiences zero translational and rotational accelerations measured along the body-fixed axes. Mathematically, this condition is represented using the following equations:

$$\dot{u} = \dot{v} = \dot{w} = 0$$
$$\dot{p} = \dot{q} = \dot{r} = 0 \quad (4)$$

where u, v, w and p, q, r are the translation velocities and body angular rates along the three coordinate directions, respectively. In steady forward flight and hover, symmetry can be assumed with respect to the body X-Z plane. Further, wing inertial loads are ignored because the inertial loads vanish when time-averaged over one flap cycle and substituting the translational and angular velocity in Eqs. (4) in (3) yields that the flight dynamics reduce to the following longitudinal trim equations:

$$\bar{F}_x = 0$$
$$\bar{F}_z = -\,mg$$
$$\bar{M}_y = 0 \quad (5)$$

The results presented in this work do not discuss trim with respect to maneuvering flight and are only limited to hover and forward flight for the sake of brevity

Trim Algorithm (Loose Coupling): While CFD can provide detailed numerical representations of the flow-field around the insect body and flapping wings, the time taken for CFD analysis is several orders of magnitude larger than the quasi-steady aerodynamic model. Using CFD directly in the trim process is therefore infeasible because of the associated computational cost. The approach adopted for this work uses CFD to sporadically correct the quasi-steady model at select stages of the trim process, instead of replacing it entirely. For this work, the delta air-loads approach used commonly in rotorcraft analysis [20,21] is adopted, schematically shown in Fig. 5. While the quasi-steady model cannot accurately predict the forces and moments for a given wing kinematic condition, it captures the trends well. Therefore, the quasi-steady model is used to compute the Jacobian (which is of negligible computational cost) and the results from CFD are used to augment the forces and moments from the quasi-steady mode

6. $F_{corrected} = F_{QS} + F_{D}$

where $F_{D} = F_{CFD} - F_{QS}$ Using this approach, the temporal resolution of the unsteady air-loads generated by CFD are judiciously introduced into the trim framework without incurring a large computational cost. The flight dynamic model with built-in quasi-steady wing aerodynamics is used to obtain the trimmed wing kinematics using Newton-Raphson iterations. For each iteration, the Jacobian (sensitivity of cycle- averaged force and moments to perturbations in the individual trim variables) is recalculated using forward differences about the current trim guess. For the vehicle parameters considered, 5–10 iterations are required until trim is achieved. As noted previously, the quasi-steady aerodynamic model is very inexpensive, and requires a few seconds to achieve trim. Note that while trim solutions are obtained using a simplified aerodynamic model with very low computational cost, three-dimensional CFD simulations are key to capturing all associated unsteady phenomena and providing fundamental insight into the nature of the flow-field.

Three flapping kinematics parameters for each wing are allowed to vary in this study to achieve the forces and moments required to maintain steady flight. These parameters are identical between the port and starboard wings in steady level forward flight, and are referred to as "trim variables" (Fig. 2)

1. Half peak-to-peak stroke amplitude ($f_{\max}$), which predominantly affects the vertical thrust.
2. Stroke flapping bias (f_{off}), which primarily affects the center of lift during the flapping stroke, and hence the body pitching moment.
3. Stroke plane angle (b), which produces a forward tilt of the resultant thrust vector, used to counteract drag by generating positive propulsive force.

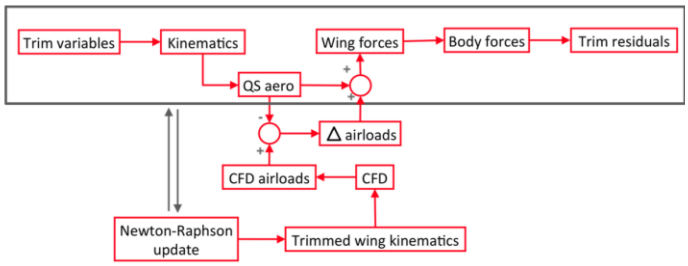


Figure 5. Schematic of the loose coupling algorithm for high-fidelity insect trim.

The trim variables, moments, and forces with respect to the global inertial (space-fixed) coordinate system are depicted in Fig. 5. Within each flap cycle, the aerodynamic forces on the wing are time-varying. Therefore, the insect body experiences non-zero accelerations when flying at constant speed and steady non-accelerating flight is obtained in an average sense over one flap cycle

Trim Controls: The variation of stroke amplitude with forward speed is shown in Fig. 6. As the forward speed increases, vehicle drag increases primarily because of the increasing wing drag. The total force required from the flapping wing increases significantly beyond 0.5 m/s, and results in a subsequent increase in the stroke amplitude. The variation of stroke phase offset (Fig. 6) affects the longitudinal bias of the center of lift ahead of the vehicle CG, shifting the mid-point of the flap cycle towards the tail for forward speed less than 1 m/s. For speeds higher than 1 m/s, an opposite trend is observed where the mid-point of the flap cycle starts to shift forward. The phase offset has a dominant effect on the pitching moment, and must be analyzed in conjunction with the stroke plane tilt (Fig. 6). The reason for the significant change in the bias angle is the increase in the stroke plane tilt for speeds larger than 1 m/s.

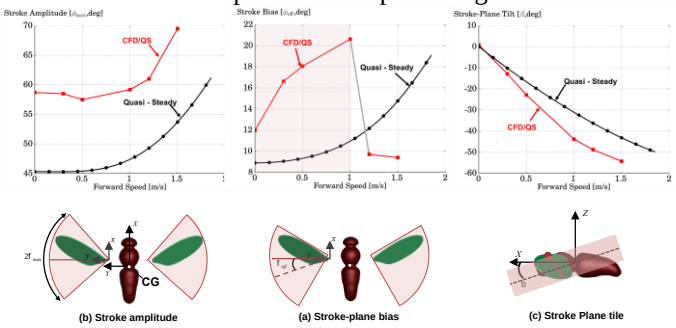


Figure 6. Variation of trim variables with forward speed: quasi-steady aerodynamics vs. CFD

Forces and Moments: Equilibrium in forces and moments in level forward flight for a speed of 1 m/s is obtained for the following trim variables, stroke amplitude ($f_{\max}$) of 59.2 deg., stroke bias (f_{off}) of 20.6 shifted from midpoint towards the trail and a stroke-plane tilt (b) equal to 43.8 tilted forward. In trimmed forward flight, the total velocity vector has components from the induced velocity, rotational velocity and freestream velocity. Due to the flapping-stroke kinematics, the wings experience different resultant velocities during the downstroke and the upstroke. The resultant LEV and circulatory forces are a function of the effective angle of attack and the resultant velocity in the flapping plane. The wings have a larger resultant velocity in downstroke than in upstroke as is illustrated in Fig. 7. This increase in dynamic pressure on the wings during the downstroke compared to the upstroke is reflected in the magnitude of the produced aerodynamic forces. Figure 7 plots the time histories of the total forces over one stroke in trimmed forward flight of 1 m/s. The total forces are contributed by the two wings (effect of insect body forces are negligible compared to those from the wings) and computed in the inertial coordinate system. As in hover, the average quantities must satisfy the trim conditions, i.e., the vertical thrust (Z) is equal to the insect MAV weight and average X and Y forces are zero. The majority of the vertical thrust is obtained during the downstroke. The downstroke and the upstroke are no longer similar due to the tilt of the plane forward and the different resulting velocity magnitude. The averaged X force is approximately zero. Most of the negative X force is obtained during the supination (with a negative peak), with a positive horizontal force obtained during the upstroke.

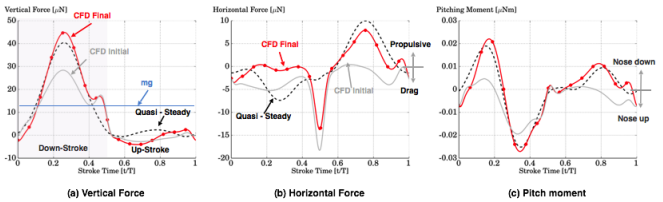


Figure 7. Time histories of aerodynamic loads during a flapping cycle in trimmed forward flight at 1 m/s: comparison of quasi-steady predictions with initial and final CFD.

B. *Performance*

The next logical step in the study of flapping wings is to evaluate its performance for a given flight condition after the insect has been trimmed, which is also an important outcome of the trim analysis. For a given vehicle weight, the mean aerodynamic power and its components can be determined at hover and different forward flight speeds. The insect mass was 1.11 mg in all cases. The presented results were obtained using CFD after converged QS-CFD loose coupling trim iterations. This reference trim condition allows for a comparison of insect flight with different wing kinematics.

Figure 8(a) presents the CFD predictions of flapping cycle averaged aerodynamic power in trimmed forward flight as a function of forward speed up to 1.5 m/s. The circles represent CFD solution and the dashed line is a least square curve-fit to better elucidate the trend. Note that the aero- dynamic power initially decreases with increasing forward speed, reaching a minimum at 0.5 m/s before increasing again for higher forward speeds. As flight speed increases, there is an increase in the flapping stroke plane tilt (β). The angle is close to zero in hover and increases to a value of 55◦ at the maximum tested forward flight speed. This trend in power is similar to those obtained from a rotor in trimmed flight. Similarities in the trend between flapping wings and edgewise rotors is apparent in a reduction in power with forward speed before increasing.

Through a CFD analysis it is challenging to isolate the components of power into induced, profile and parasitic. In the present study, power is computed through a summation of the dot product of the local force and local body velocity for all surface mesh elements. While the estimation of power is based on aerodynamic loads obtained from the CFD simulations, real insect flight (mechanical or biological) will experience losses due to inertial forces and mechanical losses in addition to aerodynamics. This aerodynamic power is decomposed into the pressure and viscous components. The physical flow around the flapping wing is highly viscous (low Reynolds number), resulting in a large boundary layer and increased shear drag, which can be referred to as profile drag. Even for low Reynolds numbers, where shear forces cannot be neglected, pressure forces still dominate shear forces.

Figure 8(b) shows the variation in shear, pressure and total power over the flapping cycle for an insect in hover. The power losses due to the pressure drag are much larger than the power losses due to the viscous profile drag through the hover stroke.

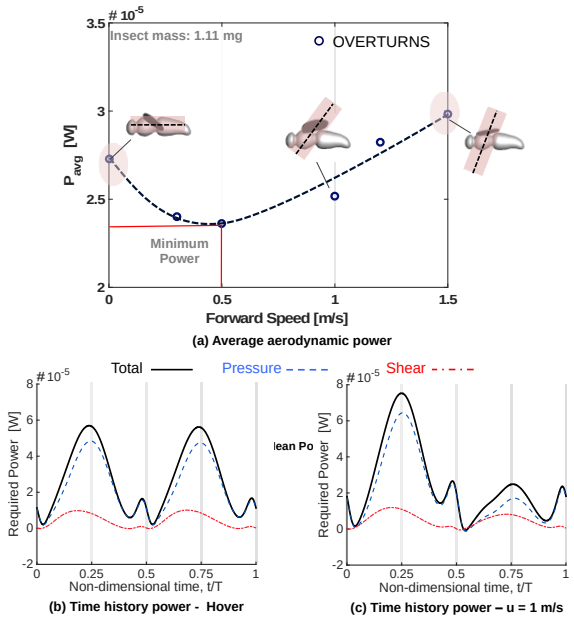


Figure 8. CFD predications of flapping wings MAV aerodynamic power mean and time history in hover and trimmed forward flight of 1 m/s.

The peak in pressure power losses occurs at the middle of the downstroke and upstroke. The viscous power losses increase at the beginning of the downstroke (and upstroke), reaching the peak around t/T = 0.2 (upstroke t/T = 0.7) and then there is a reduction until the end of the stroke, where there is a relatively small peak due to the rotational forces.

Figure 8(c) shows the time history of required power during one flapping cycle in trimmed forward flight. The larger component of total power is due to the pressure, which is related to the formation of LEV.

D. Metrics for performance evaluation: CFD capabilities developed under this task were evaluated against their validity in predicting various phenomena relevant to MAV flight. Specifically, it's ability to predict the integrated forces (lift and drag) as well as the detailed flow physics for a wide range of parameters and configurations. In addition, the CFD capabilities was combined with other disciplines: CSD to examine the benefits of flexible structures (in another CTA-MAST task), trim to investigate in more detail static stability derivatives; and will be combined to investigate free-flight response to gusty environments, and reduced order model extraction to aid in optimization. Furthermore, this work was presented in many conferences, resulting in 8 proceeding papers and one published journal article; additional journal articles will be submitted soon for publications. Furthermore, a PhD Degree was completed based on this task with the person transitioning to a post-doc position.

VI. Summary and Conclusions

An unsteady, three-dimensional, Navier–Stokes solver (OVERTURNS), was used to simulate insect flight with two flapping wings (diptera configuration). Investigated were the

aerodynamics around flapping wings, measure of efficiency metrics, analysis of trim and the performance of the insect.

The principle aim of this study was to recognize similarities between rotorcraft and flapping wings, both in terms of the results obtained and the methodology used in rotorcraft that could be transferred to the flapping wing community. Specific conclusions from this study are:

1. **Wake:** The flow field beneath flapping wings is intricate and laden with vortices, which interact with the wing resulting in complex aerodynamic phenomenon. Simplified models must account for the effects of added mass and rotational circulation along with the translational lift to accurately predict the forces and moments with varying kinematics. These effects are more predominant in forward flight as compared to hover. Flapping wings in hover show similarities to hovering rotor through the presence of axially convecting tip and root vortices. In a time-averaged sense, the wake beneath flapping wings was found to resemble the momentum theory formulation of helicopters. The contraction was more tip-domination and a loose control volume

2. **Efficiency metric:** For fixed wing, where velocity is uniform and power is the product of integrated drag and velocity, L/D is equivalent to L/P. However, this relation does not hold for flapping or rotating wings where the velocity is time-varying and non-uniform along the wing span. It can therefore be concluded that to compare between flapping wings with different kinematics or geometrical parameters as well as to other micro-air vehicles, power loading as a function of averaged lift (L/P vs. Lavg) is more appropriate as an efficiency metric than is the L/D ratio. The proposed metric properly accounts for the required aerodynamic power and vertical thrust as a function of prescribed kinematics. The ratio L/P versus lift can be used to determine the performance. Power and vertical thrust are the key parameters for design and performance. In terms of efficiency metrics, the QS model fails to predict the values obtained by CFD, especially when L/D is considered. This results from the inadequate modeling of the unsteady mechanisms in the QS model. One specific outcome of the analysis is the conclusion that for the given flapping frequencies ranging from 100-300 Hz, a pitch angle of 40 degrees has a clear advantage over pitch angles of 30 and 60 degrees in terms of L/P versus lift.

3. **Trim analysis:** A multi-fidelity approach for obtaining CFD-coupled trim solutions was adapted from rotary-wing analysis and applied to a flapping wing platform. Using a judicious number of CFD calculations to correct the quasi-steady aerodynamics so that trim is obtained with the CFD aerodynamics; thus, the detailed flowfield models can be incorporated into the analysis without incurring prohibitive run times. Wing flapping kinematics and longitudinal aerodynamic loads are cross-coupled in forward flight, i.e. one kinematic parameter affects multiple force/moment components and the sensitivity depends on the operating conditions (non-linearity). The methodology employed here considers the cross-coupling between these variables. The quasi-steady model is capable of obtaining qualitatively accurate and quantitatively representative of the integrated wing loads and the trim Jacobian (i.e., sensitivity of airloads to the wing kinematics parameters), providing a good initial condition for the coupled process. The quasi-steady model is capable of obtaining reasonable representations of the integrated wing loads and the trim Jacobian, providing both a good initial condition for the coupled process and a stable matrix for the Newton-Raphson iterations. However, this simplified model over-predicts the wing lift and under-predicts power required, and cannot capture unsteady effects during wing rotation between the up-stroke and down-stroke. CFD is therefore required for the quantitatively accurate predictions of aerodynamic power and trim kinematics that are crucial for selection of wing geometry, battery, and determination of the actuator torque required and the vehicle flight envelope restrictions.

4. **Forward flight**: The key feature in steady forward flight is that the flapping stroke plane is tilted forward, which leads to asymmetry in the forces produced and the power required during during the downstroke and the upstroke. The majority of the vertical thrust is produced during the downstroke, while the positive propulsive force is produced during the upstroke. The required aerodynamic power is larger during the downstroke, mostly due to the pressure drag resulting from flow separation presence of a LEV. The power required due to the shear forces is nearly zero during the supination and pronation because the LEV detaches and the flapping velocity slows down to zero, and non-zero during the translational portions due to the formation of a boundary layer.

The main contributions of this study to the further development of flapping wing MAVs, can be summarized as two key contributions:

1. A new efficiency metric is proposed. It is recommended evaluate designs based on the L/P ratio versus average lift to enable comparison between different platforms.

2. An alternative computationally efficient approach for trim analysis on insect-like flapping wing kinematics, has been developed and demonstrated.

4. **Power:** The required power of flapping wing in forward flight is similar to rotorcraft, in that there is a decrease from hover to a minimum power before increasing with forward flight speed. The ability to trim with CFD is key in obtaining insights into the mechanics of flight. For instance, the forces generated by the wings in forward flight are asymmetric between the up-stroke and the downstroke, with the down-stroke responsible for the vertical force and up-stroke responsible for the propulsive force. It is recognized that this division of forces in forward flight is true for a Diptera configuration, and might vary in the presence of insects with four wings.

VII. Technology Transition

The research presented briefly in this report has demonstrated the use of a CFD computational platform coupled to quasi-steady aerodynamics trim algorithm based on a loose coupling strategy. While the trim algorithm was successfully applied to obtained trimmed variables, where the system satisfy the trim conditions of moments and forces, there are several issues that need to be addressed to improve its

capabilities and its applicability to a wide range of insect flight modes. As potential ways in which these research findings may be leveraged in future work is using advanced computational technology.

The CFD-QS trim coupled algorithm is computationally efficient, when it runs on the HPC facilities. But, the computational cost can be further reduced by running the entire methodology on one or many GPUs. The current implementation uses a hole-cutting method to exchange information between multiple overset meshes that runs serially, which incurs a high computational cost. Using a different method, that can handle the data transfer in parallel, will ultimately lead to a reduction in the computational time.

REFERENCES

[1] Yang, K., Lakshminarayan, V. K., and Baeder, J. D., "Simulation of a Cycloidal Rotor System Using an Overset RANS Solver," American Helicopter Society 66th Annual Forum Proceedings, Phoenix, AZ, May 11-13, 2010.

[2] Lakshminarayan, V., and Baeder, J., "Computational Investigation of Micro-Scale Shrouded Rotor Aerodynamics in Hover", Journal of American Helicopter Society, Vol. 56, No. 4, October 2011.

[3] Malhan, R., Lakshminarayan, V., Baeder, J., Chopra, I., "Investigation of Aerodynamics of Rigid Flapping Wings for MAV Applications: CFD Validation," American Helicopter Society Specialists Conference on Unmanned Rotorcraft and Network Centric Operations, Tempe, Az, Jan 25-27, 2011.

[4] Lakshminarayan, V. K., and Baeder, J. D., "Computational Investigation of Micro Hovering Rotor Aerodynamics," Journal of the American Helicopter Society, Vol. 55, (2), April 2010.

[5] Sitaraman, J., Datta., A., Baeder, I. and Chopra, I., "Coupled CFD/CSD Prediction of Rotor Aerodynamic and Structural Dynamic Loads For Three Critical Flight Conditions", 31st European Rotorcraft Forum, Firenze, Italy, 2005.

[6] Potsdam, M., Yeo, Hyeonsoo, Johnson, Wayne, "Rotor Airloads Prediction Loose Aerodynamic/Structural Coupling", AHS Annual Forum, Baltimore, MD, 2004.

[7] Malhan, R., Lakshminarayan, V., Masarati, P., Baeder, J., Chopra, I., "Investigation of Aerodynamics of Flexible Flapping Wings for MAV Applications: Methodology Validation," 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, Honolulu, Hawaii, 23-26 April 2012.

[8] Ghosh, D., Baeder, J.D., "Compact Reconstruction Schemes with Weighted ENO Limiting for Hyperbolic Conservation Laws", SIAM Journal on Scientific Computing (accepted).

[9] Ghosh, D., Baeder, J.D., "Compact Reconstruction Weighted Essentially Non-Oscillatory Schemes for the Unsteady Navier-Stokes Equations", AIAA Applied Aerodynamics, June 25-28, 2012, New Orleans, LA.

[10] Thibault, J. C., and Senocak, I., "CUDA Implementation of a Navier-Stokes on Multi-GPU Desktop Platforms for Incompressible Flows," AIAA Aerospace Sciences Meeting, Orlando, Florida, January 2009.

[11] Corrigan, A., Camelli, F., Wallin, J., and Lohner, R., "Running Unstructured Grid Based CFD Solvers on Mod-ern Graphics Hardware," AIAA Conference on Computational Fluid Dynamics, San Anto- nio, Texas, June 2009.

[12] Chandar, C., Sitaraman, J., and Mavriplis, D., "Dynamic Overset Grid Computations for CFD Applications on Graphics Processing Units," Seventh International Conference on Computational Fluid Dynamics, Big Island, Hawaii, July 2012.

[13] Thomas, S., Amiraux, M., and Baeder, J., "GPU-accelerated FVM-RANS Hybrid Solver for Simulating Two-phase Flow beneath a Hovering Rotor," 69th Annual Forum of the American Helicopter Society, Phoenix, AZ, May 2013.

[14] Kennedy, C. A. and Carpenter, M. H., "Additive Runge-Kutta schemes for convection-diffusion-reaction equations," Appl. Numer. Math., Vol. 44, No. 1-2, 2003, pp. 139–181.

[15] Kanevsky, A., Carpenter, M. H., Gottlieb, D., and Hesthaven, J. S., "Application of implicit-explicit high order Runge- Kutta methods to discontinuous-Galerkin schemes," J. Comput. Phys., Vol. 225, No. 2, 2007, pp. 1753–1781.

[16] Faruque,I. and Humbert, J. S., "Dipteran Insect Flight Dynamics. Part1: Longitudinal Motion About Hover," The Journal of Theoretical Biology, Vol. 264, (02), 2010, pp. 538–552.

[17] MacFarlane, K. et al., "Quasi-Steady and Computational Aerodynamics Applied to Hovering Drosophila Dynamics," AIAA Applied Aerodynamics Conference, Honolulu, HI, June 27 - 30, 2011.

[18] Amiraux, M., Numerical Simulation and Validation of Helicopter Blade-Vortex Interaction Using Coupled CFD/CSD and The Levels of Aerodynamic Modeling, Doctor of philosophy, Department of Aerospace Engineering, University of Maryland, College Park, MD, 2014.

[19] Badrya, C., Amiraux, A., Baeder, D. J., and Kroningers, C. M. "Parametric Study of Insect Flapping Performance using CFD", AHS Aeromechanics Specialists' Conference, 22-24 Jan., 2014; SF, CA.

[20] Datta, A.,Sitaraman, J.,Chopra, I. and Baeder, D. J., "CFD/CSD Prediction of Rotor Vibratory Loads in High- Speed Flight," Journal of Aircraft, Vol. 43, (6), 2006, pp. 1698–1709.

[21] David, C. T., "The Relationship Between Body Angle and Flight Speed in Free-Flying Drosophila," Physiol. Entomol., Vol. 13, 1978, pp. 191–195.

Bio-Inspired Tandem-Wings

MCE-17-1.1

Kinematics, Flow Physics, Gust Tolerance and Maneuverability for MAST-Scale Aerial Platforms

Stuart Laurence, Nathan Shumway, Mateusz Gabryszuk
Department of Aerospace Engineering
University of Maryland
College Park, Maryland 20742
Email: stuartl@umd.edu

***Abstract*—Free flight tests were conducted with dragonflies in a testing environment that allows for photogrammetry to be used to measure the kinematics of dragonflies in straight and maneuvering free flight. The testing environment is designed such that the dragonfly flies through an area where a gust can be generated and the resulting response observed. Bulk kinematics are shown here for both inverted and upright straight flight. Two significant differences are seen in the kinematics: the upstroke-to-downstroke ratio changes from > 1 to < 1, and the pitch amplitude increases while the mean pitch angle decreases. When encountering a low speed lateral gust, dragonflies used asymmetric pitch, where the upwind wings had a larger pitching amplitude than the downwind wings. When performing pitching turn maneuvers dragonflies increase their flapping frequency and slightly decrease their flapping amplitude. Lastly this data also gives information on wing deformation. An example of a wing during the downstroke is shown with the same wing during the upstroke which shows that the magnitude of deformation varies significantly between the two.**

I. Introduction

MAST-scale (palm-sized) aerial platforms will require robust flight characteristics, including the capabilities for station-keeping and agile maneuvering in the presence of gusts and turbulence. Dragonflies, by virtue of their highly independent tandem-wing configuration, are remarkably agile and capable of efficient forward flight; such tandem-wing layouts are also thought to be more resistant to gusts [1]. Though dragonflies have been studied in a variety of environments there are very few studies that separate the kinematics of all four wings. Thus, in order to be able to apply the relevant principles to the design of MAST-scale platforms, a better understanding of dragonfly flight, including highly resolved wing kinematics for all four wings and detailed fluid dynamics, under both normal and gusty conditions is required.

Several studies have characterized dragonfly wing kinematics in different types of flight. For efficient cruising, the fore- and hindwings are normally out of phase; by in-phase stroking, on the other hand, dragonflies can reach speeds of 10 m/s and instantaneous accelerations of almost 4g [2,3]. Flow visualizations have shown that in counterstroking, a leading-edge vortex (LEV) forms on the forewing downstroke, with attached flow throughout on the hindwings; for in-phase stroking, the downstroke is highly separated, with a single LEV across both fore- and hindwing [4]. Preliminary computational studies have also indicated that a flapping tandem wing configuration may exhibit superior gust resistance than a single wing [1]. Characterizing realistic dragonfly wing kinematics to deduce principles useful to a MAV-designer, however, is a challenging task. The flexibility of the dragonfly wing together with its anisotropic, corrugated cross section, means that a large number of marker points and multiple camera views must be used to derive a complete description of the wing motion and deformation as it flaps [5,6].

Tracking techniques using artificial markers have been applied to dragonflies [7] and locusts [6] and is able to provide detailed information about the kinematics of dragonflies and the deformations of their wings. To contrast the response of these insects in different conditions this level of detail will be beneficial. In order to understand how dragonflies respond to gusts we use a large number of markers for each wing and constrain the dragonflies so that they fly through our visualization volume. This allows us to get both detailed and accurate data on the kinematics and deformation that dragonflies use during straight flight and when responding to gusts.

II. Related Work

In order to better understand the flow structures that dragonflies employ in straight flight and when encountering gusts, the OVERTURNS code (MCE-16-1.2) will be modified so that the kinematics and deformations measured in the experiments can be used as input. For unsteady simulations this will give detailed information about the aerodynamic characteristics of dragonflies free flight with minimal acceleration, which have to date not been fully simulated.

OVERTURNS uses a finite volume approach to solve the unsteady compressible NavierStokes equations in space and time. For spatial reconstruction of the inviscid fluxes, a third-order Leers Monotone Upstream-Centered Scheme for Conservation Laws (MUSCL) scheme with Korens [8–10] limiter can be used with an option to use higher order schemes such as the Compact-Reconstruction Weighted Essentially Non-Oscillatory (CR- WENO) schemes [11]. The viscous terms are computed using a second order central differencing scheme. Because flapping-wing simulations are performed at low Mach numbers O(10^{-3}) and low Reynolds numbers O(10 10^3), Roe-Turkel low-Mach preconditioning is applied to improve the convergence and accuracy of the compressible NS solver. This

solver has been previously used to solve for the trimmed flight conditions of a fruit-fly like model [12]; modifications are presently underway to solve for the airflow around a dragonfly in flight.

III. Approach

Experiments on free-flying dragonflies have been carried out in the AVL small-scale wind tunnel, both with and without gusts. The wing kinematics have been determined using multiple-camera photogrammetric techniques (note that as dragonflies cannot perform the clap and fling, the extent of the wing motion is more limited than, for example, in damselflies [1,13]), placing less stringent demands on the number of imaging devices. Here wing kinematics in non-gusting cases will be compared directly with those recorded under gusting conditions; these comparisons can be made between similar flight paths in both conditions.

Blue Dashers (*Pachydiplax longipennis*), as well as two other species of dragonflies, were collected at parks near the University of Maryland, with the permission of the Prince Georges County department of Parks and Recreation. Dragonflies were brought back to the university where they were placed in a refrigerator prior to markers being put on the wings using an ultra-fine point marker. As the wing sizes of different specimens varied, the pattern is not consistent between wings, but there are several common features in the different marking patterns: 6-8 marks were applied to the leading edge, 2 marks to the wing tip, and 8+ marks were applied to the trailing edge. Markers were spaced out such that they would be distinguishable by the cameras, but still give a dense enough pattern that the deformations could be measured.

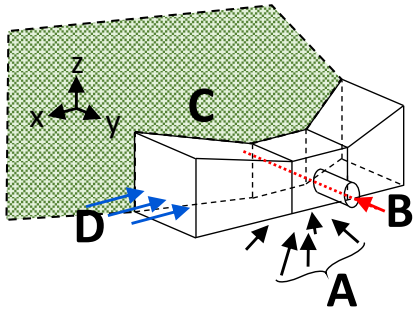


Fig. 1: Diagram of test setup. A – Arrows indicate the position and orientation of the cameras. B – The red arrow and line indicate the direction and desired flight path of the dragonfly. C – The outside area is a large netted enclosure for the dragonfly to fly into. D – The direction of the gust produced by the wind tunnel. The contraction and expansion are to connect with the AVL wind tunnel.

During initial testing of dragonflies, in the summer of 2015, dragonflies were placed in the steady flow of the wind tunnel in hopes that they would fly. Repeated testing with several different species of dragonflies made it clear that the dragonflies would not exhibit normal flight behavior as they attempted to escape from the test section. In order to increase the number of tests where we obtain usable data a new testing environment was developed. Rather than having an open flight area, a setup has been created (shown in Figure 1) where the dragonflies are forced to fly into a space where the cameras can capture the flight of the dragonfly. In order to be able to test with and without a gust of wind, the setup was designed to attach to the AVL wind tunnel, which would then provide the gust. This setup allowed cameras to be placed below the test section and spread out such that pointing angle of the cameras (indicated by the arrows in Figure 1) were at least 20° off from all of the other cameras. This angle is small for doing photogrammetric measurements, but was necessary for the cameras to all have sufficient lighting and to capture as much of the test section. The cameras used during these tests were some combination of Phantom v710s with a v310 and a v2512. All cameras were run at 7200 frames per second. At this speed all but the v310 could still operate with a resolution of 1280x800, but the v310 needed to have the resolution reduced to 768x584. The test section was illuminated by 6 100W LED flood lights.

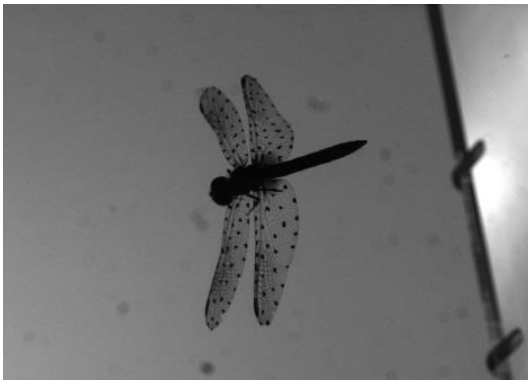
TABLE I: Total number of tests where data was acquired

Dragonfly Species	Test Type	Number of Tests
Blue Dasher	With Gust	243
	Without Gust	221
Common Pondhawk	With Gust	105
	Without Gust	68
Common Baskettail	With Gust	0
	Without Gust	32

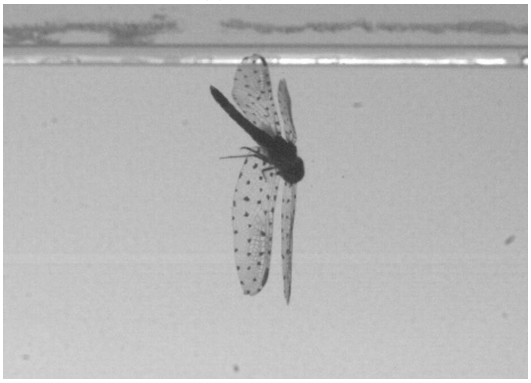
During the summer of 2016 tests were conducted with three primary species of dragonflies. All species were of similar sizes and weights as our setup required the dragonflies to be small enough to fit through the staging tube. Over the course of the testing season we were able to record videos for 669 tests. A breakdown of the types of dragonflies and conditions of the tests is shown in Table I. These videos included several different maneuvers, with climbs being the most common.

The dragonflies were inserted into a tube connected to the test section and then encouraged to fly into the test section. In most cases the dragonfly would not fly straight through the test section, but because of the constraints imposed on them by the test environment we were still able to capture usable data from about 50% of tests. Figures 2 and 3 show typical images of the dragonfly recorded during testing, as well as the level of image detail that was obtained from the cameras.

Once test data was acquired the trajectory of the dragonfly was determined by fitting a geometric model (Figure 4) to the silhouette of the dragonfly in the camera views. This was done by placing the model in a starting position and using a numerical optimization scheme to modify the position and orientation of the dragonfly to improve the fit. This was then done for every frame in the sequence to obtain the flight path and orientation of the dragonfly throughout the test.



(a) Camera 1



(b) Camera 2

Fig. 2: Sample test frame from two of the five cameras during a free-flight experiment. The cameras are angled approximately 90 degrees away from each other, allowing for increased visual coverage of the wings.

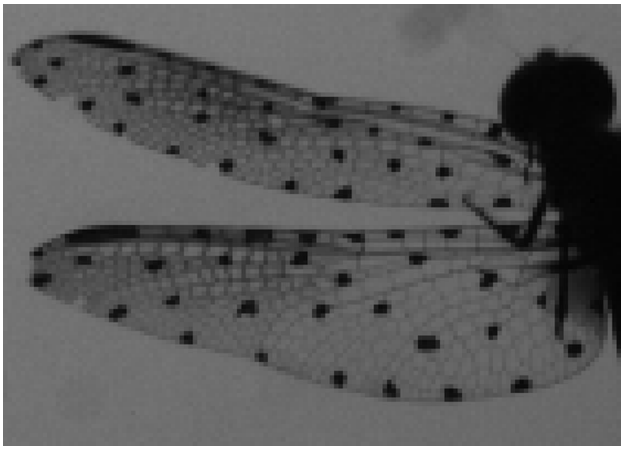


Fig. 3: Zoomed in view of dragonfly wings in 2a showing typical marker placement and the pixel-level detail of information recorded by the cameras.

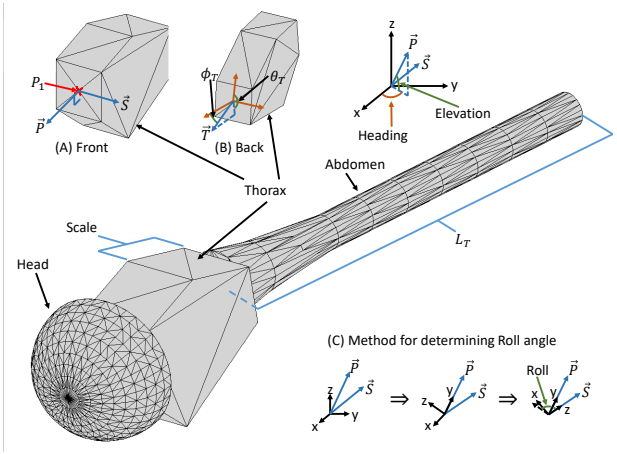


Fig. 4: Body model used to track the dragonflys position and orientation throughout the test.

Following the trajectory determination of the dragonfly, the locations of the markers on the various wings were found. There were a large number of markers on each wing (>25 on the forewings and >30 on the hindwings) but not all the markers were visible in two or more cameras at all times.

Several primary kinematic parameters were established to describe the motion of the dragonfly wings during each test sequence. These parameters are based on the trajectory of a least-squares wing plane fitted to the available marker data in each frame. The motion of the plane was tracked in a body-centered set of coordinate axes (S, P, T) as defined in Figure 4. By tracing the path of the wingtip markers throughout each sequence, stroke start- and endpoints were identified in the body frame. These were then combined with the wing root attachment point to establish individual stroke planes for each half-stroke of the wing. By comparing the stroke plane with the position of the body and wing planes, four key kinematic parameters were identified: the stroke plane inclination, stroke plane amplitude, wing flap angle, and wing pitch angle. The definitions of these parameters are shown in Figures 5 and 6.

A key step in the extraction of wing kinematics from marker position data was the calculation of the roll angle of the dragonfly in each image frame. Given that the aforementioned kinematic parameters are based on the body coordinate axes and that these axes depend on the orientation of the dragonfly, the roll angle must be determined in order to accurately calculate the kinematics of the wingstroke. The fitted body model provided an initial estimate of the roll, but the relative symmetry of the thorax about the longitudinal axis limited the accuracy of this method. Consequently, a new approach was devised to increase the accuracy of the roll angle in test frames where sufficient data was available. This was done by finding a 'wing base plane' composed of the intersection points between the leading edge lines of each wing and a cylinder with radius approximately equal to that of the dragonfly's thorax fitted to the body centerline. The orientation of this plane was then used to recalculate the roll angle of the dragonfly.

IV. Results

Of the 600+ tests where video was saved, 11 flights of Blue Dashers are presented here to give a glimpse of four different kinds of flight, straight-inverted, straight-upright, climbing,

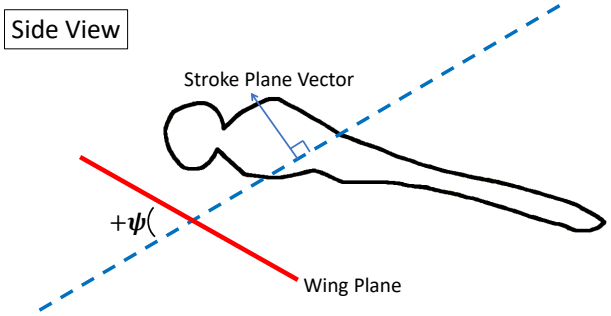


Fig. 5: Side view of dragonfly body showing stroke plane vector and wing pitch angle, ψ. The stroke plane inclination angle is the angle between the stroke plane and the body horizontal (S-P plane), while the wing pitch angle is the angle between the wing and stroke planes in the P-T plane.

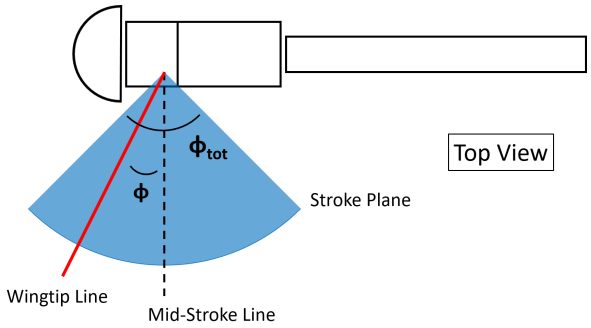
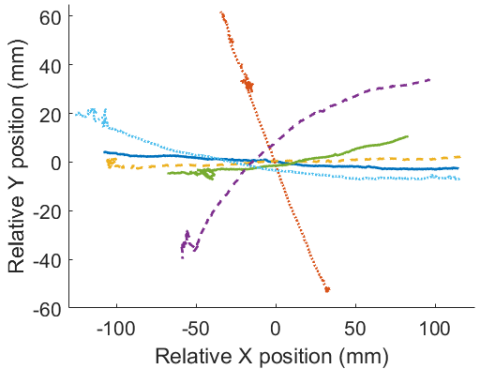


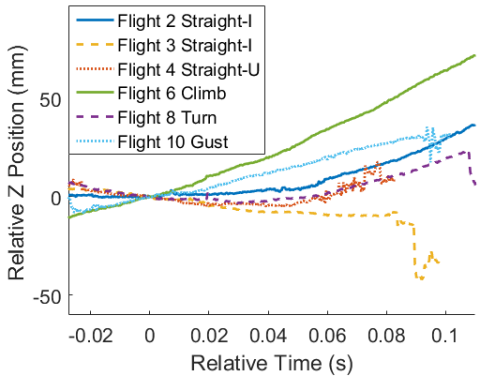
Fig. 6: Top view of dragonfly body showing stroke amplitude angle, ϕ_{tot} and wing flap angle, ϕ. The stroke amplitude angle is the angle between the stroke endpoint lines emanating from the wing root point. The flap angle is the angle between the projection of the wingtip line onto the stroke plane and the mid-stroke line of the stroke plane.

and turning, as well as straight-inverted flight while encountering a lateral gust. These tests were chosen because the data from the body model matched the various flight types and the four wings were visible in multiple cameras throughout. The preponderance of inverted flights is due to the way that the dragonflies started in the staging tube, where they would hang off of a stick until agitated and then fly directly out of the tube without significantly changing their orientation. In tests four and five the section of flight analyzed here is later in the test time and the dragonflies had righted themselves. To illustrate the various flight types the body position data for several of the tests is presented in Figure 7. All tabulated data presented in this section, such as the velocities and accelerations for the various flights in Table II, is given as the average value $\pm$ one standard deviation.

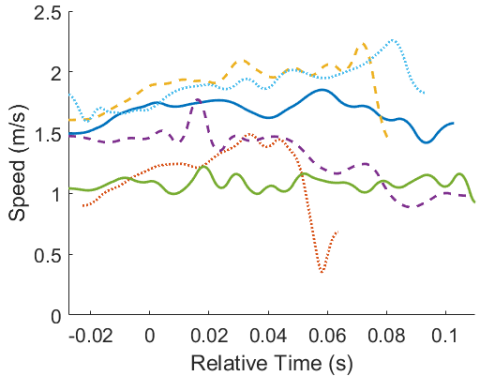
In this analysis, straight flights are flights where the x-y position data from the body model is fit well ($R^2 > 0.9$) by a linear function and the elevation change is less than 15 mm. This elevation change was thought to be acceptable since it



(a)



(b)



(c)

Fig. 7: Body position data for a sampling of the tests. (a) presents the flight path in the x-y plane, (b) shows the elevation change over the course of the flight, and (c) gives the speed of the dragonfly through the test

TABLE II: Dragonfly mass and flight information

Flight Number	Flight Type	Mass (mg)	Speed (m/s)	Acceleration (m/s)
1	Straight-Inverted	153	1.7 ± 0.1	5.5 ± 4.1
2	Straight-Inverted	220	1.7 ± 0.1	9.6 ± 8.2
3	Straight-Inverted	153	2.1 ± 0.1	5.7 ± 2.5
4	Straight-Upright	231	1.2 ± 0.1	14.6 ± 8.5
5	Straight-Upright	159	1.1 ± 0.1	8.3 ± 1.8
6	Climb	231	1.1 ± 0.1	7.5 ± 3.2
7	Climb	129	1.9 ± 0.1	8.9 ± 2.1
8	Turn	164	1.3 ± 0.2	16.9 ± 1.2
9	Turn	234	2.0 ± 0.1	22.8 ± 12.6
10	Gust	237	1.8 ± 0.1	7.3 ± 1.9
11	Gust	251	1.9 ± 0.1	8.7 ± 4.1

is a small portion of the total distance traveled (>100mm), as seen in Figures 7a and 7b. Climbing flights are flights where the dragonfly climbs at a constant rate throughout the test and the climb angle is above 10°. Turning flights were defined as flights where the position data in the x-y plane has an approximately constant radius of curvature, found by fitting a circle to the x-y position data, <300 mm. In these two tests the large accelerations are expected as the dragonfly must be sustaining 10 m/s^2 of centripetal acceleration. Finally, the gust encounters that were analyzed were picked because they closely match the straight flights, to allow for a good comparison.

The kinematic data is presented as amplitude and offset values. The flap offset value is the angle between the mid-stroke line and the plane where the wing bases are on the body model. The pitch offset value is the average pitch angle. Amplitudes are calculated from adjacent extrema in the various flights considered. The upstroke-to-donwstroke ratio are the duration of the upstroke to the duration of the downstroke. In addition to the various kinematic parameters the upstroke-to-downstroke ratio (UDR) is presented. This is the ratio of the duration of the upstroke to the duration of the downstroke. If this value is greater than one, the upstroke is longer than the downstroke and if it is less than one the reverse is true.

A. *Straight Flights*

Figure 8 shows the flap and pitch angles for the three straight, inverted flights. This data shows that there is significant variation between individual flights, making it difficult to draw sweeping conclusions from a limited data set. The flap data generally follows a sinusoidal pattern, with some variation when the dragonfly slows an individual wing. The pitch data has two different types of motion that can be categorized as peak type and plateau type motion. Peak type motion is seen in Figure 8b, where the pitch angle comes to an extrema and quickly starts changing. Plateau type motion is seen in Figure 8b where the pitch angle is held constant for the majority of the upstroke and downstroke and the rotation is confined to a period right around stroke reversal. Throughout all flights the pitching rotation happened symmetrically around the point of stroke reversal or slightly before stroke reversal.

The kinematic parameters for the straight, inverted flights are presented in Table III. For these dragonflies there is

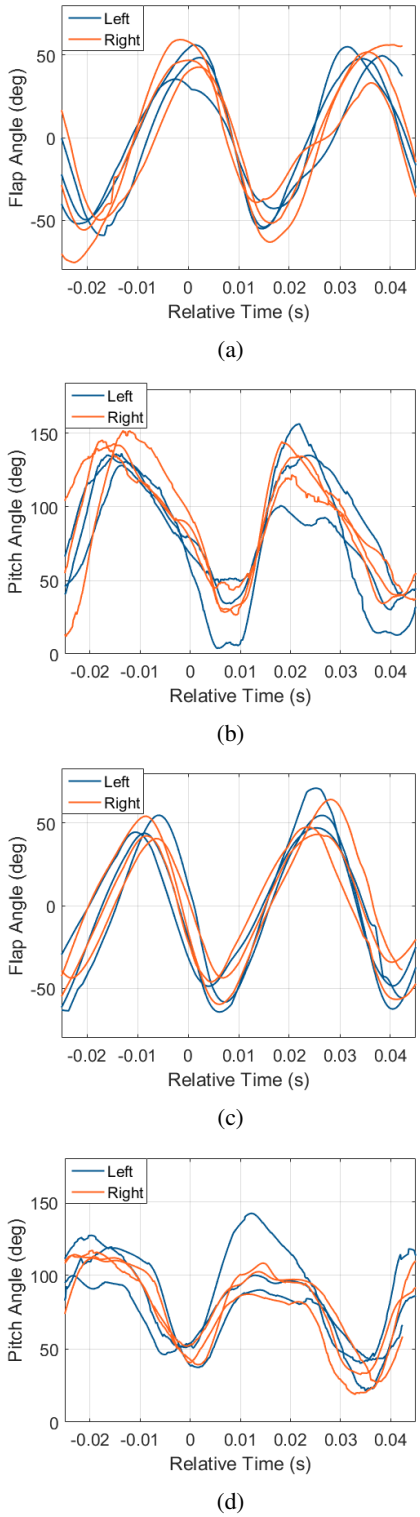


Fig. 8: Plots of wing kinematics for straight, inverted flights. (a) and (b) are plots of the forewing kinematics, and (c) and (d) are for the hindwings. t=0 is the start of a downstroke for the left hindwing in all flights.

TABLE III: Kinematic parameters for straight, inverted flight

	$\Phi(^o)$	$\phi_{off}(^o)$	$\Psi(^o)$	$\Psi_{off}(^o)$	f (Hz)	UDR	Phase $(^o)$
Left Forewing	100 ± 10	6 ± 33	99 ± 32	78 ± 11	28.7 ± 1.8	1.2 ± 0.2	0 ± 0
Left Hindwing	110 ± 12	0 ± 34	71 ± 18	77 ± 13	30.5 ± 2.1	1.2 ± 0.2	-95 ± 13
Right Forewing	101 ± 18	10 ± 31	105 ± 14	86 ± 7	28.6 ± 2.6	1.4 ± 0.3	25 ± 30
Right Hindwing	98 ± 16	-2 ± 39	65 ± 11	72 ± 12	30.8 ± 3.0	1.2 ± 0.1	-94 ± 6

minimal asymmetry between the left and the right wings. In addition, the hindwings lead the forewings by a phase angle of 90. The flapping frequency is 30 Hz, with the forewings flapping slightly faster than the hindwings. The most significant difference between the fore- and hindwing kinematics is the pitching amplitude, where the forewings have a significantly higher pitching amplitude than the hindwings.

Figure 9 shows the flap and pitch angles for both the upright and inverted straight flights. The flapping amplitude for both flight types is very similar, as is the pitching amplitude for all but the right forewing. One notable difference that can be seen in Figure 9d in is that the pitch offset for the upright flights is higher than that for the inverted flights. This difference is most significant for the hindwings, but a smaller difference is also seen in the forewings. In addition to the pitch offset angle difference, the upstroke-to-downstroke ratio for the inverted flights is above one, where it is below one for the upright flights. This means that the dragonflies are always spending more time with their wings moving in the z direction than the +z direction regardless of their orientation. There is also a notable difference in the phasing between the forewings and the hindwings, which can be clearly seen by comparing the data in Table IV and III. This is most likely due to the higher accelerations of flights 4 and 5.

B. Climbing Flights

Flights 6 and 7 represent steady climbs. The climb rates, 544 mm s^{-1} for flight 6 and 396 mm s^{-1} for flight 7, are significantly different, but the highest of the runs analyzed. Wing kinematics for the climbing flights are presented in Table 4. In these flights there is asymmetry between the left wings and the right wings in the flap angle, but that is largely from flight 7 where the dragonfly performs a roll. There is also asymmetry in the pitch angle. The wing kinematics for climbing dragonflies are not significantly different from those for straight flights.

C. Turning Flights

Two turning flights were analyzed. The dragonfly in flight 8 is turning left with a turn radius of 169 mm and the dragonfly in flight 9 is turning right with a turn radius of 248 mm. Both dragonflies were observed to perform pitching, rather than yawing, turns (where they are rolled by 90° and pitching up) and descended by 15 and 21 mm respectively. Due to marker visibility issues only one wingstroke from flight 9 was extracted, making for three total wingstrokes for analysis. A comparison between the turning flights and the straight, upright flights is shown in Figure 10. From Figure 10a and 10c it is clear that the flapping frequency for the turning flights is higher than that for the straight flight. In addition, Table VI shows that the flapping amplitude, specifically for the forewings, is lower than that of straight flight.

D. Gust Encounter

For the gust flights the dragonflies flew into a 0.3 m/s crosswind as they exited the staging tube. This is $\sim$15% of the forward flight speed of the dragonfly during the test, coming from the right side of the dragonfly. The resulting kinematics, shown in Fig. 8 and detailed in Table 5, are quite similar to those for straight, inverted flight. The most significant difference, seen in Figure 11b, is a significant asymmetry between the left and right wings in pitch. This means that the dragonfly increases the pitching of the upwind pair of wings in response to this low speed disturbance. In addition the pitching amplitudes are larger than those for straight flight.

E. Stroke Plane Orientation

The orientation of the stroke plane was analyzed for the five straight flights, two climbing flights, and two gust flights. Stroke planes were calculated for each half stroke, and are characterized by their normal vectors. As discussed here, the orientation of the stroke planes is referring to the orientation of said vector. The orientation of the stroke planes was most consistent when measured relative to the average velocity of the dragonfly, and as such the two turning flights were not included in the analysis. The resulting vectors for flights 1-5 can be seen in Figure 12.

There are two significant conclusions that can be drawn from this data. First the forewing stroke planes are pointed inward, while the hindwing stroke planes are more closely aligned with the average velocity. This is in line with the general configuration of the dragonfly's wings, which can be seen in Figure 13. The angle between the normal vector and the average velocity, in the plane from Figure 12, is $-35^o \pm 28^o$ for the left forewing, $-5^o \pm 18^o$ for the left hindwing, $22^o \pm 24^o$ for the right forewing, and $-3^o \pm 17^o$ for the right hindwing. The second conclusion is that the stroke planes are almost always pointed in the +z direction. For flights 1-5, 9, and 10 89% of the stroke plane vectors are above the horizontal plane, with an angle between the normal vector and the horizontal plane of $17^o \pm 18^o$ for the left forewing, $20^o \pm 14^o$ for the left hindwing, $21^o \pm 15^o$ for the right forewing, and $18^o \pm 14^o$ for the right hindwing. For the climbing flights these angles are even higher: $52^o \pm 19^o$ for the left forewing, $47^o \pm 28^o$ for the left hindwing, $34^o \pm 28^o$ for the right forewing, and $48^o \pm 25^o$ for the right hindwing.

TABLE IV: Kinematic parameters for straight, upright flight

	$\Phi(^o)$	$\phi_{off}(^o)$	$\Psi(^o)$	$\Psi_{off}(^o)$	f (Hz)	UDR	Phase $(^o)$
Left Forewing	93 ± 10	17 ± 29	92 ± 12	90 ± 10	31.3 ± 2.2	0.8 ± 0.1	0 ± 0
Left Hindwing	98 ± 7	18 ± 30	78 ± 11	91 ± 6	32.2 ± 0.6	0.9 ± 0.1	-72 ± 6
Right Forewing	82 ± 8	-15 ± 32	67 ± 10	92 ± 9	31.5 ± 0.9	0.9 ± 0.1	-11 ± 29
Right Hindwing	85 ± 5	-12 ± 30	65 ± 3	88 ± 3	32.0 ± 1.3	0.9 ± 0.1	-64 ± 5

TABLE V: Kinematic parameters for climbing flight

	$\Phi(^o)$	$\phi_{off}(^o)$	$\Psi(^o)$	$\Psi_{off}(^o)$	f (Hz)	UDR	Phase $(^o)$
Left Forewing	93 ± 6	-20 ± 21	121 ± 14	84 ± 3	31.8 ± 2.5	1.0 ± 0.2	0 ± 0
Left Hindwing	96 ± 12	-28 ± 29	75 ± 11	77 ± 4	33.7 ± 3.4	1.1 ± 0.3	-68 ± 16
Right Forewing	102 ± 12	29 ± 21	91 ± 12	75 ± 4	33.4 ± 4.9	1.2 ± 0.2	4 ± 26
Right Hindwing	110 ± 9	26 ± 21	82 ± 7	71 ± 15	33.2 ± 3.3	1.1 ± 0.2	-94 ± 4

TABLE VI: Kinematic parameters for turning flight

	$\Phi(^o)$	$\phi_{off}(^o)$	$\Psi(^o)$	$\Psi_{off}(^o)$	f (Hz)	UDR	Phase $(^o)$
Left Forewing	72 ± 9	17 ± 6	$99 \pm \pm 26$	100 ± 3	35.5 ± 1.1	0.9 ± 0.1	0 ± 0
Left Hindwing	85 ± 5	13 ± 11	81 ± 7	90 ± 5	33.6 ± 3.2	0.9 ± 0.1	-91 ± 4
Right Forewing	76 ± 12	-14 ± 13	97 ± 11	97 ± 8	35.4 ± 1.6	1.0 ± 0.1	-4 ± 26
Right Hindwing	90 ± 8	-18 ± 15	94 ± 9	86 ± 3	34.1 ± 3.4	0.9 ± 0.1	-79 ± 5

TABLE VII: Kinematic parameters for gust-encounter flight

	$\Phi(^o)$	$\phi_{off}(^o)$	$\Psi(^o)$	$\Psi_{off}(^o)$	f (Hz)	UDR	Phase $(^o)$
Left Forewing	103 ± 6	18 ± 20	104 ± 20	73 ± 8	29.0 ± 0.9	1.2 ± 0.1	0 ± 0
Left Hindwing	110 ± 11	8 ± 14	80 ± 14	68 ± 9	28.8 ± 0.6	1.2 ± 0.1	-84 ± 7
Right Forewing	106 ± 7	-6 ± 20	137 ± 20	100 ± 10	28.9 ± 1.5	1.1 ± 0.1	13 ± 15
Right Hindwing	91 ± 10	-17 ± 25	86 ± 25	83 ± 10	30.3 ± 1.2	1.1 ± 0.1	-70 ± 8

F. Wing Deformations

In addition to the kinematic data, wing deformation data have been obtained from Test 3. Two different times, one during the downstroke the other during the upstroke, are shown in Figure 14. The wing surface is reconstructed by fitting a third order polynomial surface to the points on the wing that were reconstructed in that frame, after removing outliers. Outliers were defined as points that were more than two standard deviations away from the mean wing plane. These indicate that the dragonfly wing bends to give the wing camber during the downstroke. During the upstroke the magnitude of the deformation is significantly higher as the wing is twisted.

V. Conclusion

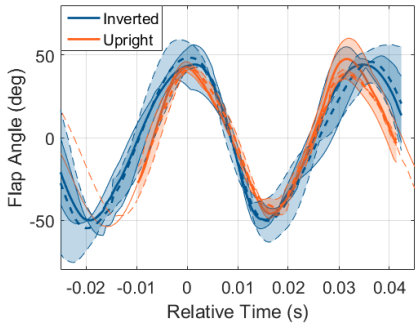
Through the implementation of photogrammetric techniques, detailed kinematic data have been successfully extracted from free-flying dragonflies in both gusting and gust-free conditions. The testing configuration allows for the study of straight free-flight, which has rarely been done, and the methodology provides a significant increase in the amount of detail obtained as compared to previous studies. The resulting data show in-phase flapping of left and right wings in gust-free forward flight, but significant phase differences between the flapping of fore- and hindwings on each side of the dragonfly.

There are several significant kinematic differences that have been observed in different flight conditions throughout this study. First there are two major differences between dragonflies flying inverted vs. upright. For inverted flights the upstroke-to-downstroke ratio is greater than one, whereas it is less than one for upright flights. This means that, in both flight conditions, the dragonfly wings spend longer moving in the -z direction than in the +z direction. In addition the pitch angle of the forewings of the inverted flights have a significantly higher amplitude where both forewings and hindwings have a slightly lower mean pitch angle. The combination of these factors allow the dragonfly to fly similar flight paths regardless of its orientation.

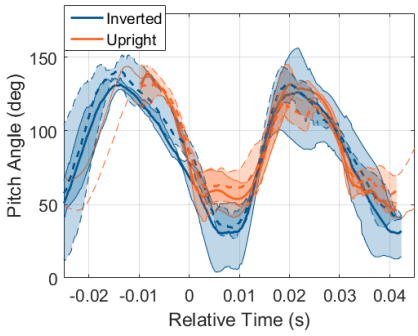
In addition to the differences between upright and inverted flight, there are significant differences between straight and turning flight. Turning flight had a noticeably higher wingbeat frequency, and slightly smaller forewing flap amplitudes. These differences are likely due to the higher accelerations required for turning flight than those seen in the straight flights.

This study represents the first measurement of dragonfly kinematics when responding to a controlled, lateral gust. The main difference between the gust response flights and the straight, inverted flights is that there is a significant left-right asymmetry in pitch angle for both the forewings and the hindwings in the gust response flights.

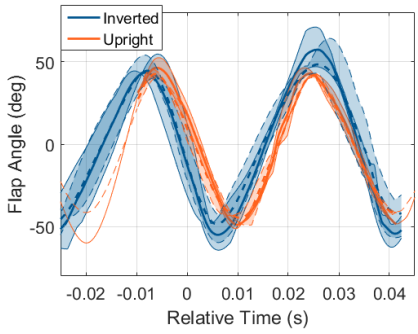
Additional research will also be continued by modifying the OVERTURNS CFD code to accept measured kinematics and deformations as an input. This will allow for accurate 3D simulations of dragonfly wings, as well as the comparison between rigid wings and realistic deformations giving insight into any aerodynamic benefits that dragonflies gain from having flexible wings. Work has already begun on developing the appropriate meshes for a Blue Dasher forewing (Figure 15). This mesh is an O-O curvilinear grid created using a 3D algebraic mesh generator [14] for structured, curvilinear, body-fitted CFD grids. This wing will be used to conduct



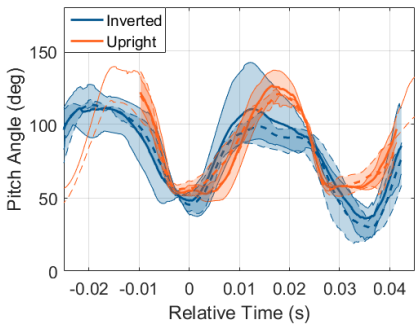
(a)



(b)

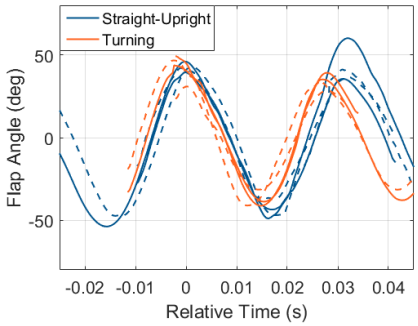


(c)

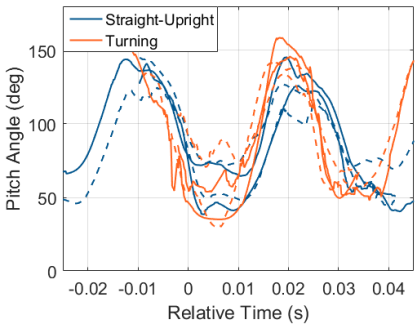


(d)

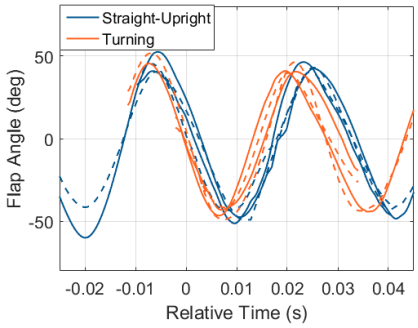
Fig. 9: Comparison of wing kinematics of upright and inverted straight flights. Thick solid lines are left wings and thick dashed lines are right wings. The shaded area is the range of data around the mean with the thin solid and dashed lines denoting the minimum and maximum values for the respective wings. (a) and (b) are forewing kinematic data and (c) and (d) are hindwing kinematic data.



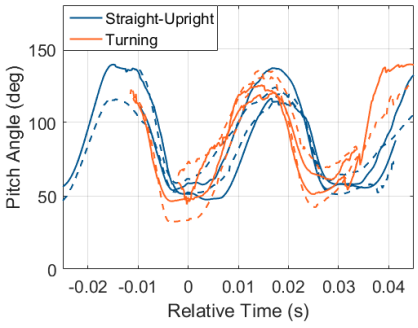
(a)



(b)

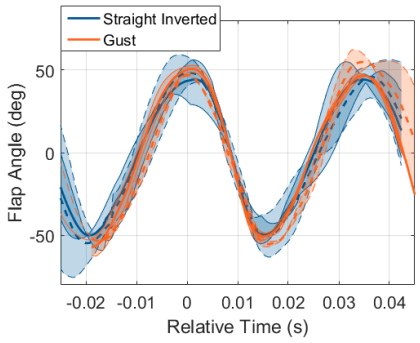


(c)

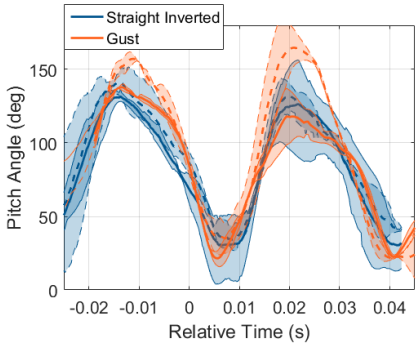


(d)

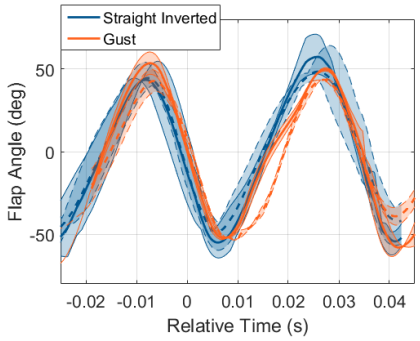
Fig. 10: Wing kinematics of turning and straight, upright flights. Left wings are designated by solid lines and right wings by dashed lines. (a) and (b) are the kinematic variables for the forewings and (c) and (d) are for the hindwings.



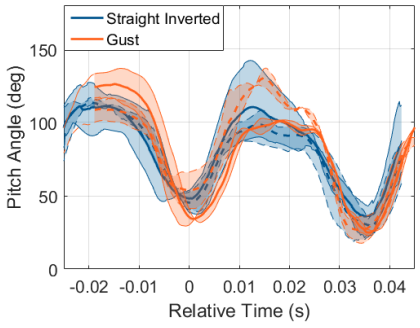
(a)



(b)



(c)



(d)

Fig. 11: Kinematics for gust encounter and straight, inverted flights. Line styles and shading follow the same conventions as Figure 9 (a) and (b) show data for the forewings while (c) and (d) show data for the hindwings.

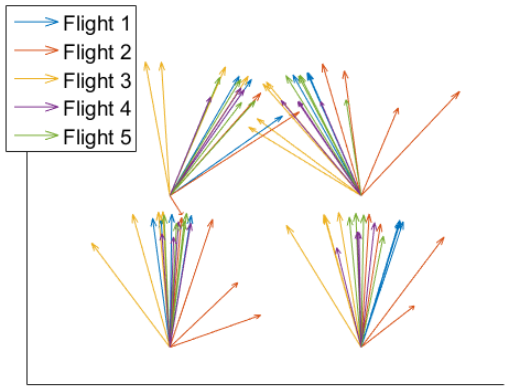


Fig. 12: Stroke plane vectors for flights 1-5. The vectors for the forewings are above those for the hindwings, and the left wings are on the left side of the plot. The average velocity of the dragonfly is straight up.



Fig. 13: Preserved Blue Dasher

simulations of a single rigid wing, then of tandem rigid wings, before moving on to the full body and deforming wings.

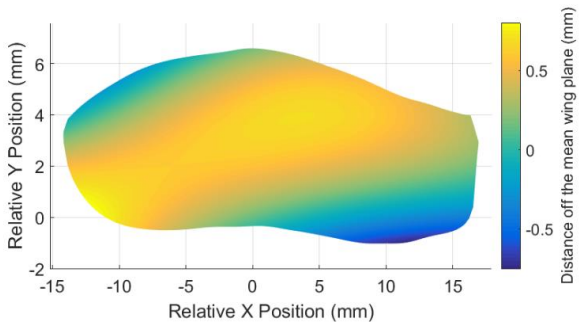
ACKNOWLEDGMENT

The authors would like to thank the Prince Georges County department of Parks and Recreation allowed dragonflies to be collected in their parks and brought to the University of Maryland. In addition camera calibration software from Simon Walker (Oxford Animal Flight Group) was used. The authors would also like to thank Sean Humbert for allowing us to use his facility and borrow his cameras and lenses. Lastly the authors would like to thank their fellow members of the High Speed Aerodynamics and Propulsion Lab for helpful suggestions in coming up with ways to analyze and present the data seen here.

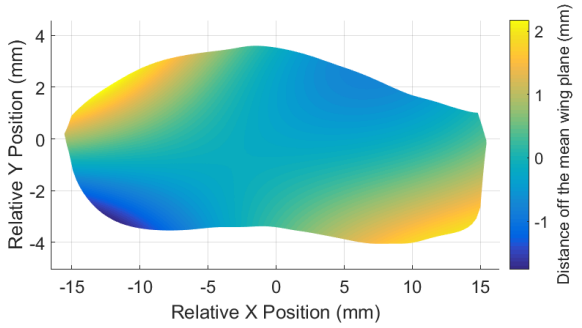
REFERENCES

[1] Y. Lian, T. Broering, K. Hord, and R. Prater, "The characterization of tandem and corrugated wings," *Progress in Aerospace Sciences*, vol. 65, pp. 41–69, 2014.

[2] D. E. Alexander, "Unusual phase relationships between the forewings and hindwings in flying dragonflies," *Journal of Experimental Biology*, vol. 109, no. 1, pp. 379–383, 1984.



(a)



(b)

Fig. 14: Two instances of dragonfly wings deformed during flight: 14a is near the end of the downstroke and 14b is near the middle of the upstroke

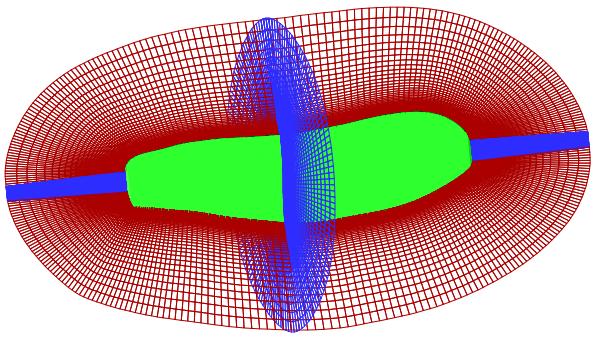


Fig. 15: A mesh for a Blue Dasher forewing to be used with the OVERTURNS CFD code

[3] M. L. May, "Dragonfly flight: power requirements at high speed and acceleration," *Journal of experimental biology*, vol. 158, no. 1, pp. 325–342, 1991.

[4] A. L. Thomas, G. K. Taylor, R. B. Srygley, R. L. Nudds, and R. J. Bomphrey, "Dragonfly flight: free-flight and tethered flow visualizations reveal a diverse array of unsteady lift-generating mechanisms, controlled primarily via angle of attack," *Journal of Experimental Biology*, vol. 207, no. 24, pp. 4299–4323, 2004.

[5] J. Wakeling and C. Ellington, "Dragonfly flight. ii. velocities, accelerations and kinematics of flapping flight." *Journal of experimental biology*, vol. 200, no. 3, pp. 557–582, 1997.

[6] S. M. Walker, A. L. Thomas, and G. K. Taylor, "Photogrammetric reconstruction of high-resolution surface topographies and deformable wing kinematics of tethered locusts and free-flying hoverflies," *Journal of The Royal Society Interface*, vol. 6, no. 33, pp. 351–366, 2009.

[7] C. Koehler, Z. Liang, Z. Gaston, H. Wan, and H. Dong, "3d reconstruction and analysis of wing deformation in free-flying dragonflies," *The Journal of experimental biology*, vol. 215, no. 17, pp. 3018–3027, 2012.

[8] K. Yang, V. K. Lakshminarayan, and J. D. Baeder, "Simulation of a cycloidal rotor system using an overset rans solver," in *at American Helicopter Society 66th Annual Forum Proceedings, Phoenix, AZ*, 2010.

[9] V. K. Lakshminarayan and J. D. Baeder, "Computational investigation of microscale shrouded rotor aerodynamics in hover," *Journal of the American Helicopter Society*, vol. 56, no. 4, pp. 1–15, 2011.

[10] R. Malhan, V. K. Lakshminarayan, J. Baeder, and I. Chopra, "Investigation of aerodynamics of rigid flapping wings for mav applications: Cfd validation," in *AHS Specialists Conference*, 2011, pp. 25–27.

[11] D. Ghosh and J. D. Baeder, "Compact reconstruction schemes with weighted eno limiting for hyperbolic conservation laws," *SIAM Journal on Scientific Computing*, vol. 34, no. 3, pp. A1678–A1706, 2012.

[12] C. Badrya, K. MacFarlane, J. Baeder, and J. S. Humbert, "Insect kinematics in trimmed flight at low reynolds numbers using cfd," *53rd AIAA Aerospace Sciences Meeting*, January 2015, aIAA Paper 2015-0045.

[13] Y. Chen, M. Skote, Y. Zhao, and W. Huang, "Dragonfly (sympetrum flaveolum) flight: Kinematic measurement and modelling," *Journal of Fluids and Structures*, vol. 40, pp. 115–126, 2013.

[14] M. Amiraux, "Numerical simulation and validation of helicopter blade-vortex interaction using coupled cfd/csd and the levels of aerodynamic modeling," Ph.D. dissertation, University of Maryland, 2014.

Robust Comprehensive CFD/CSD Methodology for Multidisciplinary Design Optimization of MAST-Scale Aerial Vehicles

Task Number: MCE-17-1.4

James Lankford, David Mayo, Ria Malhan, Inderjit Chopra
Department of Aerospace Engineering
University of Maryland – College Park
College Park, MD, U.S.A.
lankford@terpmail.umd.edu, chopra@umd.edu

Pierangelo Masarati
Dipartimento di Ingegneria Aerospaziale
Politecnico di Milano
Milano, Italy
pierangelo.masarati@polimi.it

***Abstract*—The objective of this study was to develop and validate a state-of-the-art, coupled aeroelastic analysis capable of reliably predicting the performance and detailed flowfield surrounding a highly flexible flapping wing. In conjunction with computational studies, detailed experimental tests were performed to provide data against which to validate the aeroelastic solver as well as improve understanding of flapping wing flight. Experimental studies include particle image velocimetry flowfield tests, aerodynamic force measurements and dynamic wing deformation motion tracking. Overall, the coupled aeroelastic solver results showed good agreement with the results of the aforementioned experimental tests. The final goal will be to utilize the aeroelastic analysis to carry out systematic parametric studies to come up with an optimized geometry and kinematics for efficient hover and forward flight performance of flapping-wing MAVs at MAST scale.**

***Index Terms*—micro air vehicles, flapping wings, particle image velocimetry, motion tracking, aeroelastic analysis**

I. Introduction

A. Background and Technical Challenges

Within the Army's MURI program (2003-2009) and CTA-MAST program (2008-2014, task MCE-1.4), research was focused towards developing and validating a coupled computational structural dynamics (CSD) and computational fluid dynamics (CFD) methodology to predict the aeromechanic phenomena of MAST-scale aerial vehicles. Both rotary-wing and flapping-wing based platforms were considered. Even though significant progress was made in this very ambitious aeroelastic prediction methodology, significant barriers were encountered, including inadequate controlled and documented test data at small scales, computational simulation complexities associated with low Reynolds number aerodynamics and non-existing nonlinear 2-D structural modeling. The goal of this task was to further refine this coupled multi-body CFD/CSD methodology, systematically expand validation studies for a range of realistic MAST-scale configurations and develop reduced order models for multidisciplinary design optimization. A robust comprehensive methodology will not only drastically reduce the enormous effort involved with experimental trial and error platform designs, but will also help to discover revolutionary configurations for specific Army missions.

In the early phase of this endeavor, we developed a finite-element based aeroelastic analysis for both flapping-wings and rotary-wings. The flapping-wing was modeled as a thin plate undergoing bending and torsion, and unsteady aerodynamic forces were based on Wagner and Kussner indicial functions in conjunction with Palhamus leading-edge vortex theory [1]. Limited validation studies were carried out with in-house experimental test data obtained using a biomimetic flapping-pitching mechanism. It became quite apparent from this that a linear analysis with simplified aerodynamics was inadequate to capture the aeromechanic phenomena of flapping-wings [2], [3]. This study highlighted the need to incorporate nonlinear 2-D structural models and high-fidelity unsteady aerodynamic models, as pertaining to the low Reynolds number regime, in order to accurately simulate the fluid-structure interactions and their significant impact on the force production of a flapping wing.

In the subsequent phase, an in-house CFD solver (OVERTURNS) was coupled with a topological-independent multi-body structural dynamics solver (MBDyn) to simulate the aeromechanics of flapping-wings [4]–[7]. For this, the CFD solver was extensively modified to model the highly vortical flows on flapping-wings in the low Reynolds number regime. Initially, MBDyn could model only 1-D beams undergoing nonlinear extension, bending and twisting. We actively collaborated with Politecnico di Milano to modify the multi-body solver to include 2-D shell structures.

It is noteworthy that many of the existing experimental and computational studies associated with flapping-wings

were quite limited. Most studies incorporated rigid wings, instead of elastic wings, undergoing simplistic flapping kinematics, while other studies idealized the wing structure with 1-D beams instead of 2-D plates/membranes. Often the wing models were too simplistic, undergoing low-frequency flap motion and adopted linear structural deformation instead of nonlinear analysis [8]–[15]. Again, it became apparent that the idealistic constraints presented in these studies do not provide an accurate representation of the aerodynamics and structural deformations associated with realistic flexible flapping-wings. In-house studies showed that flightworthy flapping-wing concepts require highly flexible, lightweight wings, which need to operate at high flapping frequencies and amplitudes. Once the coupled CFD/CSD prediction capability was developed, the next barrier for validation was the scarcity of detailed documented test data for flapping-wings within a controlled environment.

Thus, considerable effort was focused on generating in-house test data on low aspect ratio flapping-wings for a wide range of controlled test conditions in a wind tunnel [16], [17]. Realistic rigid and flexible wings were built and detailed flow components, using particle image velocimetry (PIV), and loads, using a microbalance, were measured for a range of flapping frequencies and amplitudes. This test data was used to validate our prediction methodology and guide us to further refine our CFD/CSD simulations. The simulation was able to capture the formation of the salient aerodynamic features produced by the flapping wing at different chordwise locations. These aerodynamic features include the leading-edge vortex (LEV), the trailing edge vortex, and small eddies that form a turbulent shear layer which leaves the wing's trailing edge (i.e. the shed wake) as it moves through its stroke. The studies showed that a 50% larger vertical lift (F_Z) and a 42% larger propulsive thrust (F_X) is generated by the flexible wing as compared to the rigid wing demonstrating that wing flexibility plays a very important role in enhancing aerodynamic force production [17].

With the current prediction capability, we have now started understanding the performance gains in flapping-wings due to structural flexibility. It is clear that flexible wings are capable of simultaneously producing positive lifting force and positive propulsive thrust over a given flap cycle in forward flight conditions. This is an important insight into how dynamic twist due to structural flexibility can provide a means of passive temporal and spanwise pitch modulation for a flapping wing MAV design. It is necessary to understand how aeroelastically tailoring of a given wing design can create a more optimal twist distribution along the wingspan throughout the flap cycle and thus improve aerodynamic efficiency and performance. This will be vital for the design and development of next-generation highly maneuverable MAVs

B. Research Objectives

The objective of this research was to further develop the aforementioned CFD/CSD prediction methodology. Improvements to the computational aeroelastic solver include the addition of advances in structural modeling and CFD mesh deformation schemes. Validation studies were expanded to cover both rigid and flexible wings in hover and forward flight. For validation, detailed experimental tests were carried out in a wind tunnel where the flowfield (using PIV) and airloads (using microbalance) were carefully recorded and documented. In addition, insect-based kinematics were systematically addressed. Flowfield measurements (using PIV) and passive wing deformations (using VICON motion tracking) were conducted and compared against coupled CFD/CSD predictions. Parametric studies were carried out to obtain insights in the aeromechanics of flapping wings under different flight conditions and assess the influence of wing flexibility on aerodynamic performance. Once fully developed, this methodology could be a paradigm shift in the flapping-wing MAV design process, which is currently performed in an ad hoc, trial and error manner.

C. Technical Approach Summary

The in-house coupled CFD/CSD aeroelastic solver was subject to a number of enhancements. Additionally, experimental studies of MAST-scale components were expanded to further validate the computational analysis. Below is a summary of the key areas that were addressed.

a) Laminated Plate Element: Structural modeling was expanded to include large deflections of 2-D twisted laminated plate element and account for bend-twist and extension-twist composite couplings. The effect of structural tailoring on flapping performance and figure of merit was examined systematically.
b) Refined CFD Modeling: Improvements to turbulence modeling (DES/DDES) were included in the CFD simulation and its effect on LEV, vorticity distribution, flow separation and induced power were examined for different flight conditions. CFD mesh deformation was made more robust to expand the range of possible wing designs to analyze.
c) Wing Kinematics: Complex insect-based wing kinematics were included and its effect on aerodynamic efficiency were examined systematically. Predicted and measured wing kinematics (using VICON) will be compared.
d) Solution Processing: In order to improve the computational efficiency of the aeroelastic prediction methodology and design optimization formulations, algorithms were parallelized to take advantage of multiple processor capabilities.
e) Experimental Testing: Detailed experimental tests were carried out in controlled environments in a wind tunnel where the flowfield (using PIV), loads (using microbalance), structural properties and motions (using VICON) were documented carefully. A range of simplified flapping-wings were built and tested using high-frequency insect-based kinematics. Test data was used for validation studies. Additionally, several new experimental test datasets are being created and documented.

II. Related Work

The work presented for this task is a continuation of the work performed in task MCE-1.4 titled *Multi-Body Microsystem Analysis Code for Rotary-Wing and Flapping Wing Systems* under CTA-MAST from 2008 to 2014. External collaborations include work with Dr. Pierangelo Masarati from Politecnico di Milano. Collaboration efforts include the further development of the multi-body structural dynamics code (MBDyn) which was coupled with the in-house computational fluid dynamics solver (OVERTURNS). Currently, the focus has shifted from studying avian-based flap kinematics in forward flight to insect-based flap kinematics in hover.

III. Technical Approach

A. Experimental Test Wings

Experimental tests were conducted on two sets of MAV-scale flexible flapping wings to provide quality test data against which to validate the predictive capability of the coupled CFD/CSD aeroelastic analysis. The first test wing has a rectangular planform with a wing span of 12.7 cm and a chord length of 7.62 cm resulting in an aspect ratio of 1.67. The wing leading edge is made of a solid titanium rod with a diameter of 1.5875 mm. Wing ribs are made of a single layer of unidirectional carbon fiber. Each rib has a width of 6.80 mm and a thickness of 0.30 mm. A thin, heat-resistant film is used to cover the wing and act as the wing membrane. A four-bar mechanism was designed to provide a sinusoidal, one degree-of-freedom flap motion. Figure 1 provides an image of the test wing.

The design of the second test wing is inspired by wings used on the flapping-wing MAV developed by Coleman et al. [18]. However, for the tests conducted, the wing was uniformly scaled, increasing the dimensions of the wing geometry by 50%. The test wing is composed of two carbon fiber spars which run along the wing leading edge (leading edge spar) and across the wing plane (cross spar). The leading edge spar has a diameter of 2.0 mm and the cross spar has a diameter of 0.5 mm. A third spar, composed of plastic shim stock, supports the wing root (root spar) and has a length of 7.70 cm, a width of 5.0 mm and a thickness of 0.762 mm. The wing frame is covered with a foam mesh fabric which has a thickness of 1.0 mm. This lightweight, flexible, low aspect ratio wing has a length of 19.05 cm, a mean chord length of 5.26 cm, a wing area of 100.13 cm^2 and an aspect ratio of 3.62. The wing has a bilinear taper with the taper ratio transition point at 14.71 cm along the length of the wing. The chord length values at the wing root, transition point and wing tip are 7.70 cm, 4.55 cm and 0.28 cm respectively. A photo of the MAV flapping wing can be seen in Fig. 2.

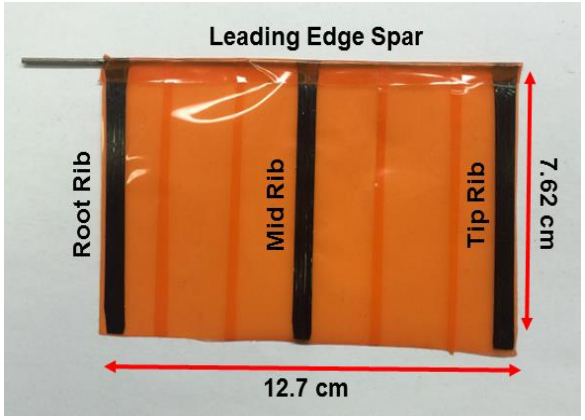


Fig. 1. Rectangular flapping wing with film covering

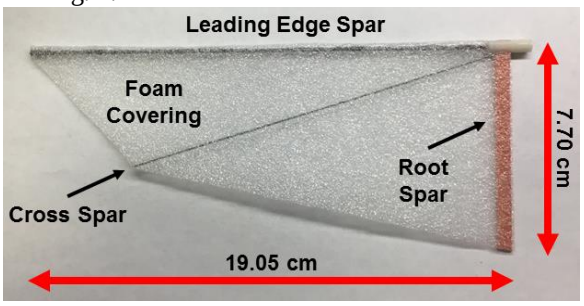


Fig. 2. MAV flapping wing with foam covering

B. Particle Image Velocimetry Setup

Time-resolved, two component PIV measurements using a double-pulsed Nd-YLF Litron LDY304 laser were performed to quantify the flowfield about the rectangular flapping wing. For all PIV tests, the images were recorded with a 4 Mega Pixel Phantom v311 time-resolved PIV camera. Forward flight tests were conducted in an open-circuit subsonic wind tunnel with a 50.8 cm by 71.12 cm test section. The air speed was kept constant at 3.0 m/s. The wing was undergoing avian-like flap kinematics and flowfield measurements were recorded at mid-downstroke and mid-upstroke between 30% and 90% span. Flap frequency was varied from 4 Hz to 10 Hz.

Hover tests were conducted in quiescent flow (U_∞ = 0.0 m/s). The wing was undergoing insect-like flap kinematics with flap frequencies varying from 4 Hz to 10 Hz. Flowfield measurements were recorded at 25%, 50% and 75% span at various instances of the flap cycle. Table 1 and Table 2 show the experimental parameters used during the flexible flapping wing experiments in forward flight and hover respectively.

TABLE 1. PIV experimental test parameters (forward flight)

Parameter	Value
Freestream Velocity, U_∞	3 m/s
Reynolds number, *Re*	15,000
Flap amplitude	±40°
Flap frequency, *Hz*	4 – 10
Reduced frequency, *k*	0.32 – 0.80
Initial pitch angle, θ_o	0° – 24°, $\Delta\phi$ = 3°

TABLE 2. PIV experimental test parameters (hover flight)

Parameter	Value
Reynolds number, *Re*	14,200 – 42,600
Flap amplitude	±40°
Flap frequency, *Hz*	4 – 10
Reduced frequency, *k*	0.341
Flap angle, ϕ	8° – 64°, $\Delta\phi$ = 8°
Span Location	25%, 50%, 75%

C. VICON Motion Tracking Setup

A VICON motion capture system was utilized to determine the structural properties of the various wing components for both the rectangular and MAV flapping wings. Additionally, the VICON system was used to measure the passive wing deformations, due to inertial and aerodynamic loading, over one complete flap cycle for the hover tests. The motion capture test setup consisted of six VICON T40-S cameras operating at 100 frames per second (fps) when performing static wing deformation tests to calculate the wing structural stiffness. Note that the structural properties for the experimental rectangular wing and MAV wing used in the hover tests can be found in Table 3 and 4 respectively. For the dynamic flapping wing tests, the six cameras were run at 400 fps to accurately capture the wing deformation over time. Reflective markers were placed at the leading edge, mid-chord and trailing edge of each rib in order to track the motion of those particular locations on the rectangular wing. On the MAV flapping wing, reflective markers were placed along the chord line at 0%, 25%, 50% and 75% span as well as a marker placed at the wing tip.

TABLE 3. Rectangular experimental wing structural properties (hover)

Property	Value
Wing length, *b*	12.7 cm
Wing chord, *c*	7.62 cm
Membrane thickness, *h*	0.06 mm
$\mathbf{EI}_{\text{root rib}}$	1.7142e-4 N-m^2
$\mathbf{EI}_{\text{mid rib}}$	2.2316e-4 N-m^2
$\mathbf{EI}_{\text{tip rib}}$	2.7014e-4 N-m^2
$\mathbf{EI}_{\text{spar}}$	2.9104e-2 N-m^2
$\mathbf{GJ}_{\text{spar}}$	0.016 N-m^2
$\mathbf{E}_{\text{membrane}}$	4.016e9 N-m^{-2}
$\mathbf{\nu}_{\text{membrane}}$	0.45
Rib mass	0.1456 grams
Spar mass	1.1360 grams
Membrane mass	1.03 grams

TABLE 4. MAV experimental wing structural properties (hover)

Property	Value
Wing span, *b*	19.05 cm
Mean wing chord, $\bar{c}$	5.26 cm
Membrane thickness, *h*	1.0 mm
$\mathbf{EI}_{\text{root spar}}$	3.47e-4 N-m^2
$\mathbf{GJ}_{\text{root spar}}$	4.96e-5 N-m^2
$\mathbf{EI}_{\text{cross spar}}$	3.27e-4 N-m^2
$\mathbf{GJ}_{\text{cross spar}}$	2.42e-4 N-m^2
$\mathbf{EI}_{\text{leading edge spar}}$	4.33e-2 N-m^2
$\mathbf{GJ}_{\text{leading edge spar}}$	3.21e-2 N-m^2
$\mathbf{E}_{\text{foam sheet}}$	1.0e7 N-m^{-2}
$\mathbf{\nu}_{\text{foam sheet}}$	0.42
$\mathbf{\rho}_{\text{carbon fiber}}$	1433.0 kg/m^3
$\mathbf{\rho}_{\text{plastic}}$	1505.0 kg/m^3
$\mathbf{\rho}_{\text{foam sheet}}$	30.0 kg/m3

D. Force Measurement Setup

Figure 3 provides a picture showing the MAV flapping wing experimental setup. All tests were carried out at a constant flap stroke amplitude of 80° at various flap frequencies. Since the tests were conducted in hover, this results in a reduced frequency of 0.156 for all cases. The forces produced by the wing during actuation were measured using an ATI Nano-17. The ATI Nano-17 is a miniature six-component balance with a resolution of ±0.318 gram-force and a resonant frequency of 7200 Hz. The force balance is mounted at the wing root to capture the total forces and moments produced by the wing. The flap angle at a given instance of the flap cycle is measured using a US Digital MA3 miniature absolute magnetic shaft encoder with a resolution of ±0.7°. The measurements from the Nano-17 are output to an IFPS box which is connected to a National Instruments USB DAQ system (NI USB-6251). The voltage signals from the force transducer are recorded and post-processed in MATLAB using a calibration matrix provided by the manufacturer to obtain the six forces and moments.

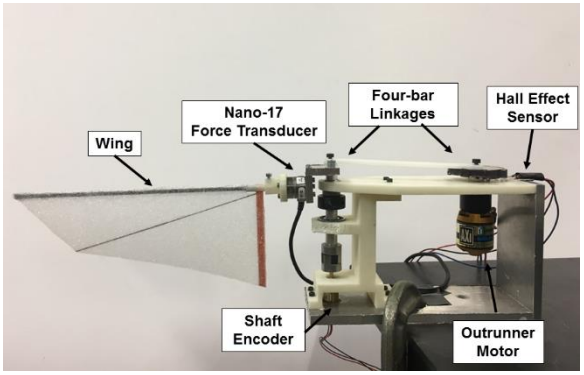


Fig. 3. MAV flapping wing experimental setup

E. CFD/CSD Model

A C-O type mesh was utilized with 267×141×85 nodes in the wrap-around, spanwise and normal directions respectively to represent the rectangular experimental test wing. Similarly, a C-O mesh type with 227×91×93 nodes in the wrap-around, spanwise and normal directions was used to represent the MAV flapping wing. The wing planforms of both CFD meshes were created to match their respective experimental test wings. A higher concentration of grid points were placed near the leading edge, trailing edge and wing tip to better capture the highly vortical flow expected near those regions. An image of the CFD meshes can be seen in Fig. 4 and Fig. 5. To capture the highly vortical flow in the wake, a Cartesian background mesh was implemented with 127×195×162 nodes in the x, y and z directions respectively. Node points are highly clustered within the initial 2.5 chord lengths from the wing surface in all directions to improve spatial resolution of the shed wake. The farfield boundary conditions are applied 10 chord lengths from the wing surface.

The structural model used for the rectangular flapping wing is made up of 46 nodes with 20 beam elements (representing the wing ribs and leading edge spar) and 32 shell elements (representing the thin wing film) in a 4×8 grid. The MAV flapping wing is made up of 85 nodes with 23 beam elements (representing leading edge, cross and root spars) and 66 shell elements making up the wing membrane.

A computational framework, written in Python, was employed to exchange data between the CFD solver (OVERTURNS) and the CSD solver (MBDyn). The structural deformations and aerodynamic loads are exchanged at the fluid-structure interface (i.e. wing surface) using a non-conformal mapping scheme, which preserves the work between the two domains.

Data transfer between the CFD and CSD solver is performed in a tight manner meaning that structural deformation and aerodynamic loading data is passed at each time step of the coupled solver solution. A more detailed description of the wing design, experimental setups and computational aeroelastic model can be found in [19,20].

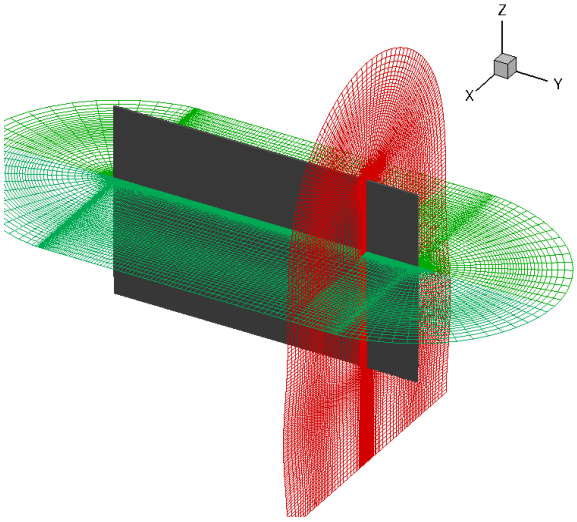


Fig. 4. C-O CFD mesh of rectangular flexible flapping wing

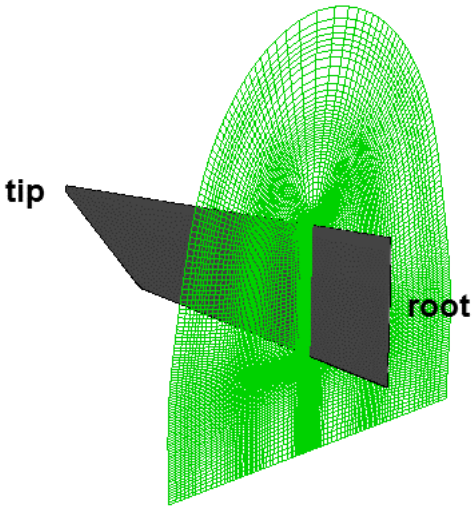


Fig. 5. C-O CFD mesh of MAV flexible flapping wing

IV. Results

A. Rectangular Flapping Wing – Forward Flight

1) Qualitative and Quantitative Flowfield Comparison

For the forward flight tests performed on the flexible flapping wing, PIV flowfield measurements as well as time-averaged force balance airloads were measured and compared to the coupled aeroelastic analysis results. Given the highly unsteady nature of the aerodynamic regime about the wing, the vorticity field is used to compare the measured and predicted flowfields. Figure 6 shows the vorticity contours and normalized velocity profiles for the 4 Hz flap frequency case for an initial root pitch angle of θ_o = 12° at 50% span.

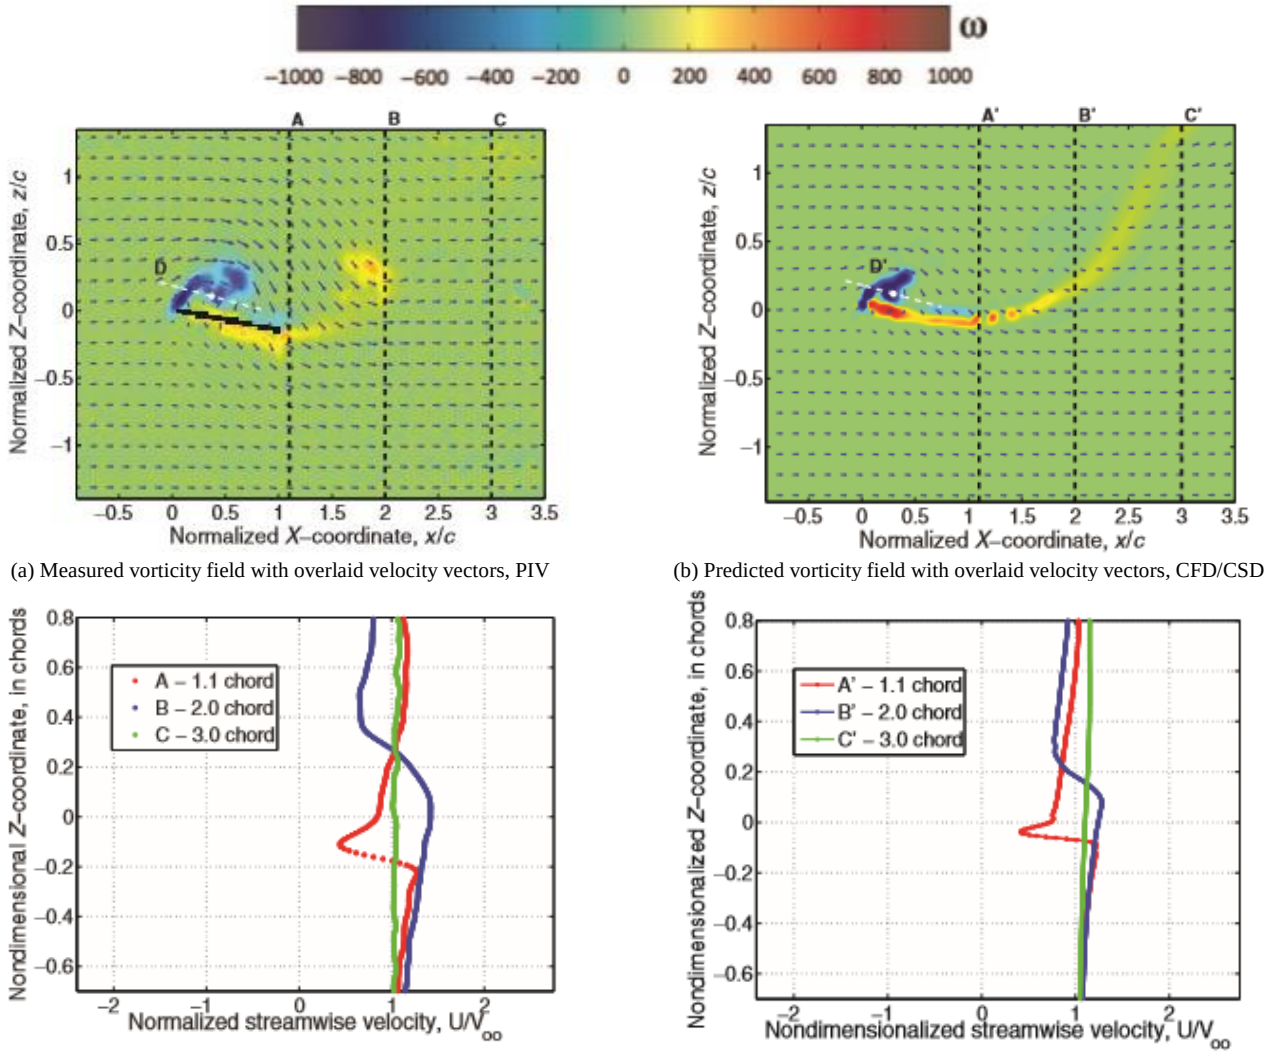
Comparing the vorticity fields in Fig. 6(a) and Fig. 6(b), we see that the predicted results exhibit good qualitative agreement with the measured results. In both cases, the leading edge vortex is seen to be attached to the wing and a notable shear layer has formed from the wing's trailing edge. While the coupled CFD/CSD simulation is able to predict many of the salient features, there are some slight discrepancies in the trajectory and size of the shed vortices (especially at the trailing edge) when comparing the experimental and computational results.

To perform a quantitative comparison, downstream velocity cuts were taken along the z-direction at x/c = 1.1, 2.0 and 3.0 where x/c = 0 is the leading edge. The downstream velocity cuts have been labeled by the letters A, B, C and A', B', C' for the PIV results (Fig. 6c) and CFD/CSD results (Fig. 6d) respectively. The predicted velocity profiles show good agreement against the measured values. The simulation is able to quantitatively capture the variation of velocity within the wake including the influence of the trailing edge shear layer. Additional validation results can be seen in [17].

2) Rigid and Flexible Wing Force Comparison

The results of the coupled aeroelastic analysis were utilized to better understand the flow physics and aerodynamic force generated by flexible flapping wings in forward flight. More specifically, the goal was to analyze the influence of wing compliance on the production of lift and propulsive thrust. Figure 7 shows the variation in vertical force coefficient (C_Z) and horizontal force coefficient (C_X) over the course of one flap cycle for the flexible and an ideally rigid wing. The results shown are for a flap frequency of 6 Hz and initial root pitch angle of 12°. Note that positive values of C_X designate a drag producing force while negative values of C_X represent a propulsive thrust producing force. Additionally, horizontal lines are displayed on the plots in Fig. 7 to show the time-averaged C_Z and C_X.

In Fig. 7(a), it is observed that a majority of the positive vertical force production occurs during the downstroke. The positive initial root pitch angle creates an asymmetry in vertical force production resulting in a positive net vertical force. However, there is a significant difference in the net horizontal force produced between the rigid and flexible



(a) Measured vorticity field with overlaid velocity vectors, PIV

(b) Predicted vorticity field with overlaid velocity vectors, CFD/CSD

(c) Measured downstream wake velocity profiles, PIV

(d) Predicted downstream wake velocity profiles, CFD/CSD

Fig. 6. Vorticity field and wake velocity profile comparison of PIV and CFD/CSD flowfields at Ω = 4 Hz, θ_o = 12°, mid-downstroke

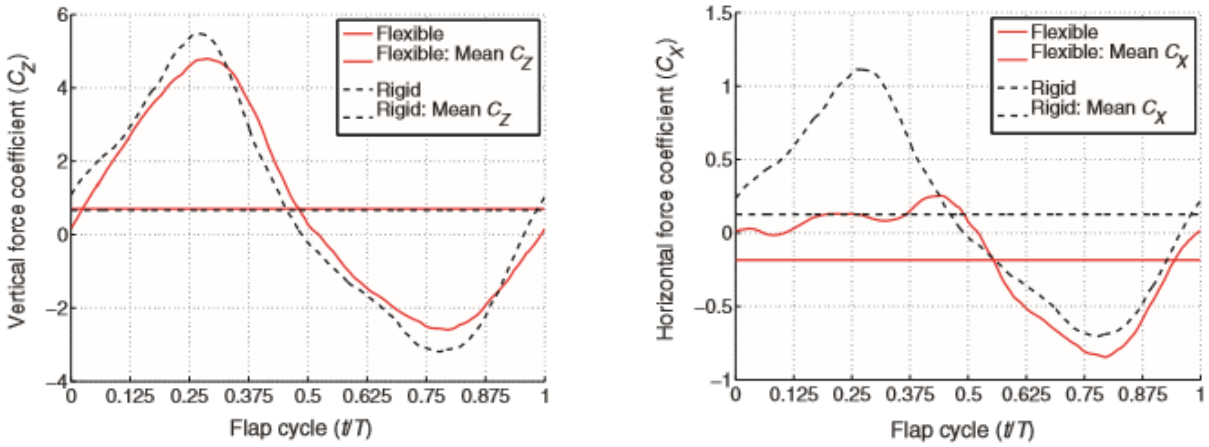
wings. In Fig. 7(b), during the upstroke, both the rigid and flexible wings produce a majority of their propulsive thrust. During the downstroke, the flexible wing produces a significantly lower amount of drag in comparison to the rigid wing. This allows for the flexible wing to produce a net propulsive force while the rigid wing produces a net propulsive drag. From Fig. 7, we see that introducing flexibility allows for the wing to simultaneously produce net lifting force and net propulsive force over one flap cycle.

B. Rectangular Flapping Wing – Hover

1) Wing Deformation Comparison

For the tests conducted on the flexible flapping wing in hover, in addition to the flowfield comparisons, efforts were made to compare the predicted passive wing deformations to those measured experimentally. Figure 8 shows the variation of the wing deformation at the trailing edge from the experimental and computational results at Ω = 4 Hz (Fig. 8(a)) and Ω = 8 Hz (Fig. 8(b)).

Comparisons are made at the trailing edge of the root, mid and tip ribs over the course of a flap cycle. Flap cycle time is non-dimensionalized by the flap period (*T*). The curves displaying the experimental results are an average of 40 flap cycles while computational results are taken from the last flap cycle of the simulation. For the 4 Hz flapping case (Fig. 8(a)), the computational results compare well against the experimental data with respect to peak displacement magnitude and phasing. However, there is some higher harmonic content displayed in the experimental data that is not captured in the computational results. In Fig. 8(b) (Ω = 8 Hz), similar to the 4 Hz case, the computational results exhibit good agreement with regard to peak displacement magnitude, especially for the mid and tip ribs. The experimental results exhibit a phase shift with the peak magnitude in wing displacement occurring after stroke reversal (stroke reversal occurring at t/T = 0.0, 1.0 and t/T = 0.5). The computational results predict this phase shift but not to the same extent as seen in the experimental data. Additionally, there is a significant level of higher harmonic



(a) Vertical force coefficient time history (b) Horizontal force coefficient time history

Fig. 7. Comparison of instantaneous force production, Ω = 6 Hz, θ_o = 12°

content (notably for the root rib) that is not predicted by the computational analysis.

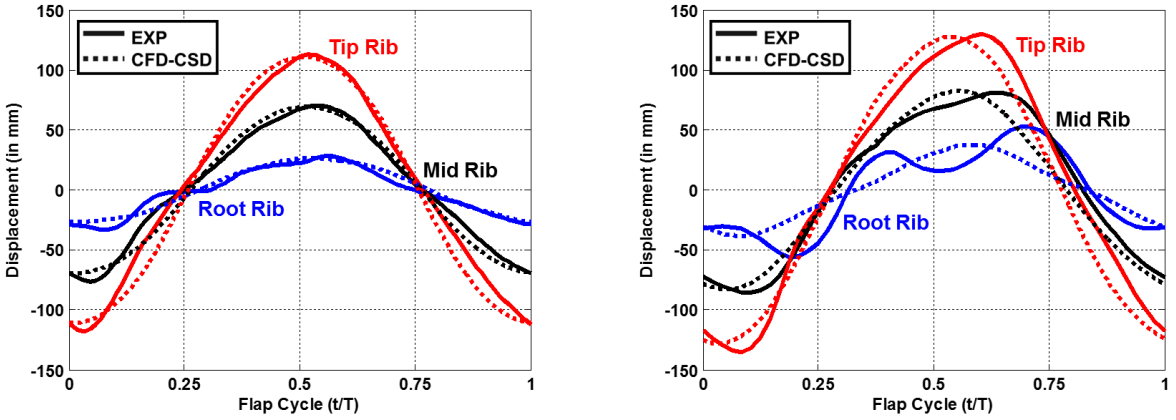
In general, the coupled CFD/CSD analysis is capable of predicting the passive wing deformations at lower flap frequencies. However, at higher flap frequencies there are noticeable differences between the measured and predicted wing deflection with respect to phasing and higher harmonic oscillations. Additional validation results can be seen in [18].

2) Effect of Flexibility on Aerodynamic Force Production

During experimentation, only a single flexible wing was tested (structural properties of the wing can be found in Table 3). The coupled CFD/CSD aeroelastic analysis is used to assess the effect of chordwise flexibility on wing performance. For the 4 Hz flapping frequency case, computational cases were carried out at 75%, 50%, 25% and 10% of the initial flexural stiffness (EI) for each of the ribs described in Table 3. Figure 9 shows the instantaneous lift (Fig. 9(a)) and drag force (Fig. 9(b)) variation, respectively, for the aforementioned varied stiffness cases over the course of one complete flap cycle. As seen in Fig. 9(a), decreasing chordwise stiffness increases lift force coefficient magnitude and has a significant effect on the variation of instantaneous lift force over time. However, in Fig. 9(b), decreasing chordwise flexibility seem to have little impact on the drag variation over time and peak drag coefficient magnitude for all the cases tested.

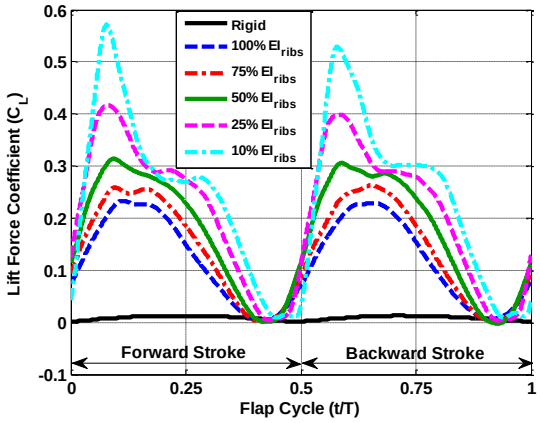
In Fig. 9(a), for stiffness values greater than 50% of the initial rib flexural rigidity, the lift coefficient time history tends to follow a periodic, nearly sinusoidal variation similar to the sinusoidal variation of the flap angle. But, as stiffness is decreased further, the lift variation begins to deviate from a sinusoidal variation. For the 25% and 10% chordwise stiffness cases, the C_L variation exhibits a large, initial spike after stroke reversal (approximately t/T = 0.08 and t/T = 0.58) followed by a plateau in C_L magnitude and finally begins to decrease after passing midstroke (t/T = 0.25 and t/T = 0.75).

Figure 10 displays the mean C_L and mean C_D over a flap cycle as a function of stiffness. The x-axis in Fig. 10 represents the stiffness for a specific case (EI) as a ratio of the initial chordwise stiffness (EI_o). From Fig. 10(a), we see that mean C_L varies monotonically and exhibits a linear trend when wing compliance is present, decreasing in magnitude as chordwise stiffness is increased. The mean C_L for the 10%

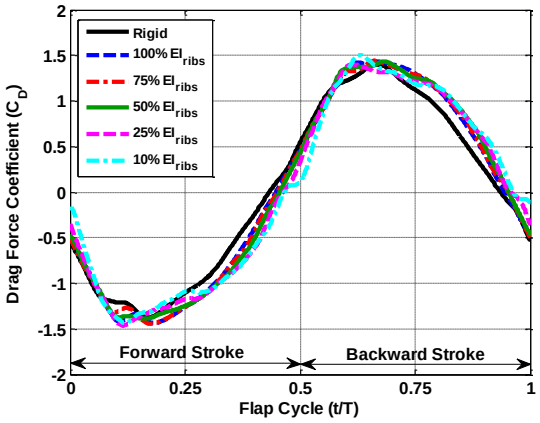


(a) Ω = 4 Hz (b) Ω = 8 Hz

Fig. 8. Comparison of experimental and computational passive wing deformations

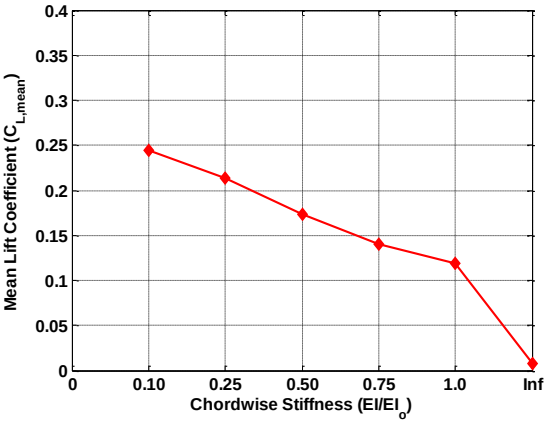


(a) Lift force-time history

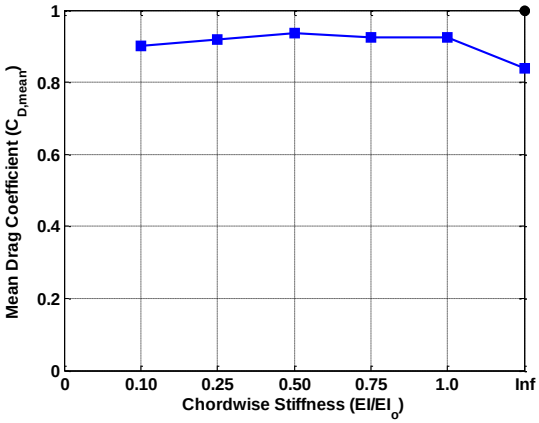


(b) Drag force-time history

Fig. 9. Instantaneous force-time history over a flap cycle, varied chordwise stiffness



(a) Mean lift coefficient variation



(b) Mean drag coefficient variation

Fig. 10. Time-averaged lift and drag coefficient versus chordwise stiffness

chordwise stiffness case is roughly twice that of the 100% stiffness case. For the flapping kinematics considered, some flexibility in the wing is necessary to produce net lift as the "infinitely rigid" case produced negligible lifting force. In Fig. 10(b), there is only a slight variation in mean C_D with chordwise stiffness. However, it should be noted that there is a noticeable increase in net drag for all the flexible wing cases in comparison to the "infinitely rigid" flapping wing case.

3) Effect of Flexibility on Wing Deformation

Figure 11 shows the variation of wing pitch at the root, mid and tip rib locations for the 100% and 10% chordwise stiffness cases over the course of a flap cycle. Pitch angle (θ) is defined as the angle between the wing flapping plane and a line joining the leading and trailing edge points at a given spanwise station of the wing. The wing is set at an initial pitch angle of θ_o = 90°. The largest wing deflections occur after stroke reversal resulting in a minimum wing pitch of θ = 82° and θ = 70° for the 100% and 10% chordwise stiffness cases respectively. The variation in wing pitch for the two cases is similar to the variation in instantaneous lift coefficient for the respective cases in Fig. 9. For the 100% chordwise stiffness case, the variation in wing pitch along the wing is nearly sinusoidal and the 10% chordwise stiffness case exhibits a spike in pitch magnitude followed by a plateau. From the results presented, temporal and spatial wing pitch variation has a strong impact on lift force variation.

This suggests that aeroelastic tailoring of a given wing design can be used to improve overall aerodynamic force production and increase wing performance.

C. MAV Flapping Wing – Hover

1) Aerodynamic Force-time History Comparison

Figure 12 shows the measured and predicted instantaneous lift and drag force variation for the MAV flexible flapping wing at a flap frequency of 6 Hz. The experimental lift and drag force-time histories are an average of 10 separate trials. Each trial is an average of the force variation over 40 flap cycles. The shaded region bordering the experimental force-time histories represent the standard deviation of the 10 averaged trials. The predicted lift and drag force-time histories are the results of the last flap cycle of the coupled aeroelastic simulation. Note that lift force is reported in grams and positive lift is said to act along the

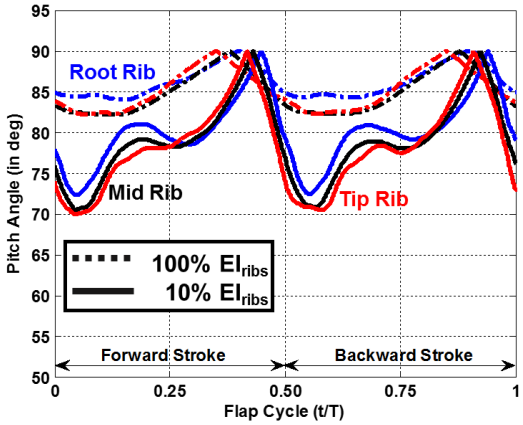


Fig. 11. Comparison of wing pitch variation for the 100% and 10% chordwise stiffness cases

positive Z-axis. Drag is also reported in grams and acts in opposition to the wing motion.

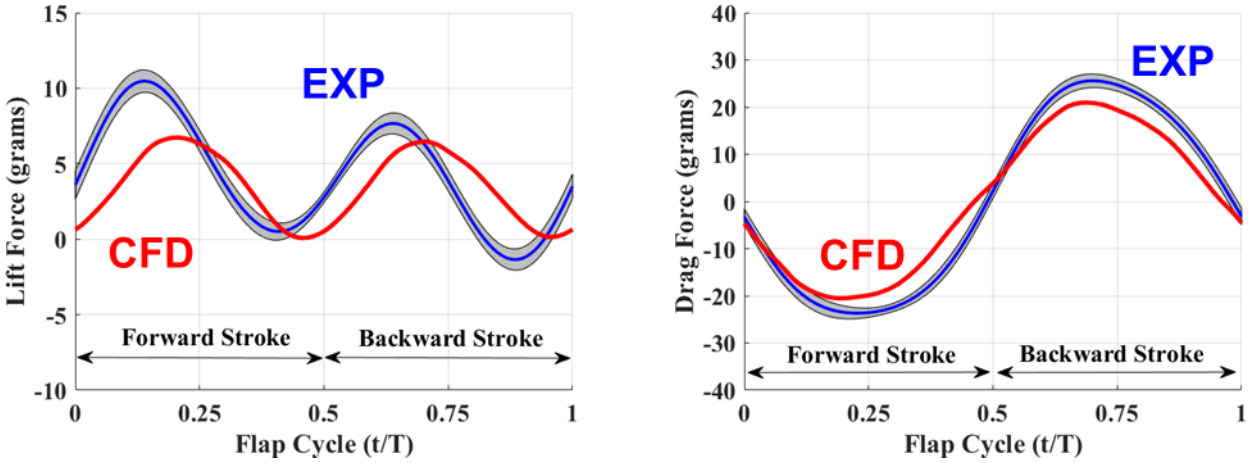
In Fig. 12(a), we see that for a majority of the flap cycle, the experimental lift force-time history is positive with some negative lift being generated near the end of the backward stroke. There is a distinct asymmetry in the lift force variation between the forward and backward strokes. The forward stroke exhibits a higher peak lift force in comparison to that of the backward stroke. Overall, standard deviation in the instantaneous lift force-time history is highest at the peaks in lift and lowest near stroke reversal (t/T = 0.5 and 1.0). When comparing the measured and predicted results, we see that the coupled aeroelastic analysis is under-predicting the peak lift force generated by the wing. Additionally, the predicted results exhibit a much more symmetrical variation in lift force between the forward and backward strokes with lift being mostly positive throughout the flap cycle. In Fig. 12(b), the experimental and CFD/CSD drag force-time histories show similar variations over the flap cycle. Similar to Fig. 12(a), the coupled aeroelastic analysis under-predicts the peak drag value generated by the wing.

2) Influence of Flexibility on Wing Performance

The results in the following sections are for a set of computational studies in which the flexural stiffness (EI) of the root and cross spars (as shown in Table 4) are varied keeping all other wing properties constant. Figure 13 shows the variation in the instantaneous lift (Fig. 13(a)) and drag (Fig. 13(b)) coefficients over a flap cycle for the 6 Hz flap frequency at varied wing stiffness values. In Fig. 13(a), decreasing the stiffness of the root and cross spars causes an increase in the peak lift coefficient during the forward and backward strokes. For the initial 100% EI case, the peak lift coefficient value is 0.266 and for the 25% EI case the peak lift coefficient value increases by 15.8% to 0.308. Comparing the lift variation for the 50% EI and 25% stiffness cases, the peak lift attained is nearly equal for both cases. However, after peak lift is produced, the 25% EI case maintains a larger magnitude of lift force especially near mid-stroke (t/T = 0.25 and 0.75). Analyzing Fig. 13(b), there seems to be little variation in the drag force-time history with respect to wing stiffness. There is a slight phase shift in the drag force variation when wing flexibility is present compared to an infinitely stiff wing. However, the peak drag coefficients are similar between all cases.

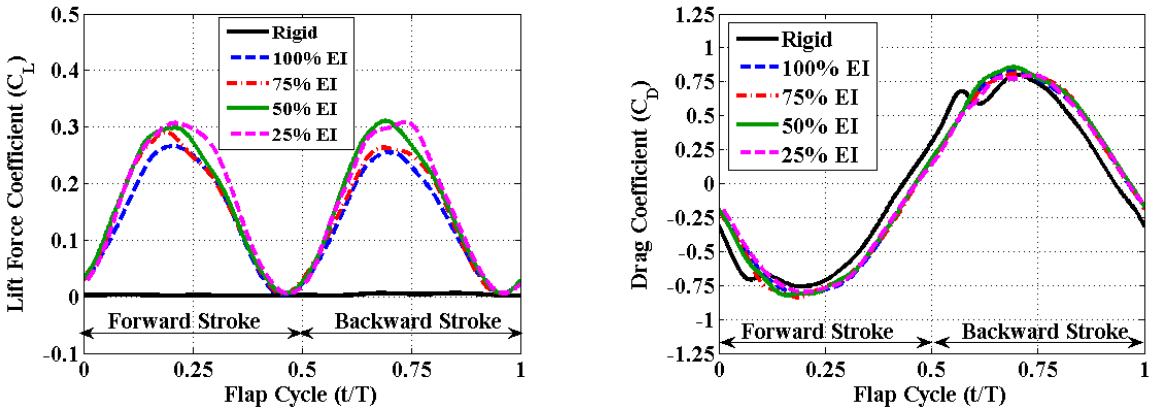
The plots in Fig. 13(a) provide a view of how wing flexibility influences the instantaneous lift generated by the entire wing over a flap cycle. However, varying wing flexibility influences the spanwise variation of wing deflection throughout the flap cycle. Thus, it is important to assess the sectional variation in lift force over the flap stroke to gauge how varied wing compliance impacts temporal and spatial production of lift. Figure 14 is a contour plot of the variation in sectional lift coefficient over the flap plane during the forward stroke.

Fig. 14(a) and Fig. 14(b) are the sectional lift contour plots for the 100% EI and 25% EI cases respectively. For both stiffness cases, the largest sectional lift coefficient values are produced during the first half of the forward stroke. Toward mid-stroke, the sectional lift produced at the wing root and wing tip diminishes. This is most likely due to the evolution of the root and tip vortices which interfere with



(a) Lift Force-time History (b) Drag Force-time History

Fig. 12. Comparison of Experimental and CFD/CSD Aerodynamic Force-time Histories



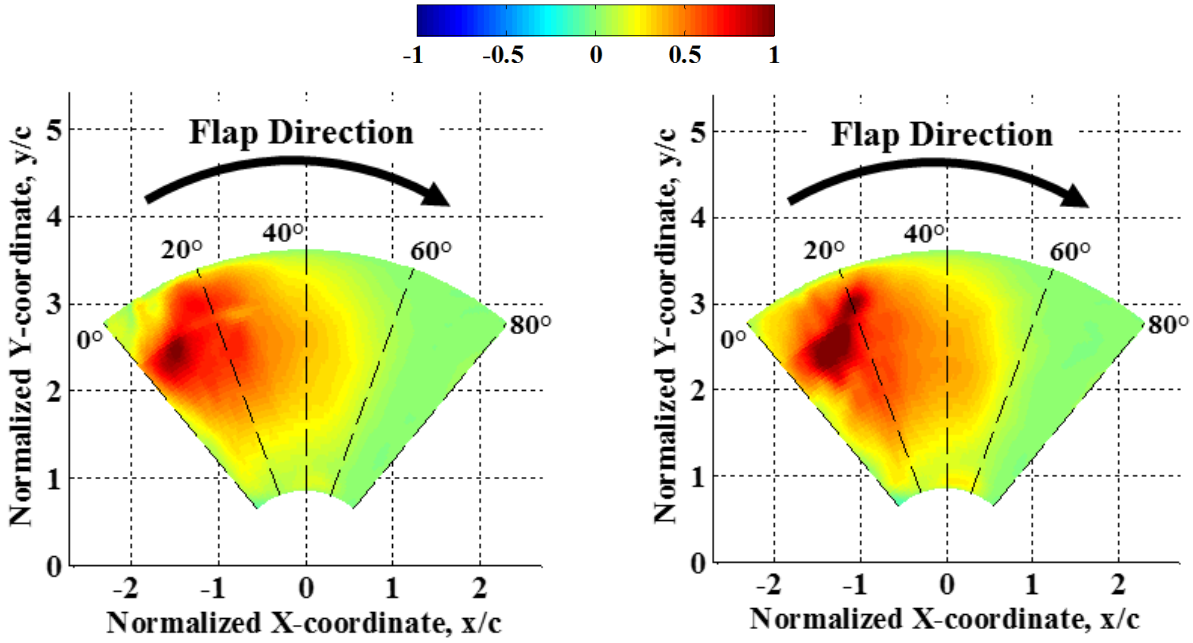
(a) Lift coefficient vs flap cycle (b) Drag coefficient vs flap cycle

Fig. 13. Instantaneous aerodynamic force coefficient variation, Ω = 6 Hz, varied stiffness

the development of the leading edge vortex and consequently decrease lift production. The 25% EI case produces higher sectional C_L values for a greater portion of the flap stroke, especially near the wing tip. Additionally, the magnitude of sectional lift is slightly greater near the wing root for the 25% EI case compared to the 100% EI case.

V. Conclusions/Future Work

The goal of the present studies was to validate the predictive capability of a coupled computational fluid dynamics/computational structural dynamics (CFD/CSD) solver to calculate the unsteady aerodynamic mechanisms due to rigid and flexible wings undergoing avian-like and insect-like flapping kinematics at MAV-scale Reynolds numbers. Particle image velocimetry (PIV) experiments were conducted to investigate the flowfield in the immediate vicinity of the wing's surface at different span locations and instances of the flap cycle. VICON motion tracking experiments were performed to measure wing deformation throughout a flap cycle. Additional force measurement studies were performed on a wing modeled after an existing flapping-MAV wing. The model analyzed by the coupled aeroelastic solver was constructed to replicate the kinematics, geometry and structure of the wings used during experimentation. Results from the experiments were utilized in validating the coupled CFD/CSD analysis by comparing the aerodynamic vorticity fields, wing trailing edge deflections over time and force-time histories. The validated CFD/CSD model was further used to understand the impact of wing flexibility on the unsteady aerodynamic forces.



(a) Sectional C_L contour, 100% EI (b) Sectional C_L contour, 10% EI

Fig. 14. Contour of sectional lift force produced during the forward stroke, Ω = 6 Hz

Future studies will involve investigating more complex wing structures and geometries that resemble those used in realistic flapping-wing micro air vehicles. Currently there is still a lack of detailed force, flowfield and wing deformation data on highly flexible flapping wings with well characterized structural properties. Additionally, the coupled aeroelastic solver will be used to further explore the design space of possible flapping wing configurations. Parametric variations from a baseline wing configuration will be assessed in order to provide trends on how design variables such as wing geometry, flapping kinematics and structural tailoring influence its performance. The goal will be to determine an optimal combination of wing parameters that will enable maximum performance in hover and forward flight by improving the aerodynamic efficiency and minimizing the power required. The aeroelastic analysis will provide insight into the extent of practical flapping-wing configurations within the design space.

REFERENCES

[1] E. C. Polhamus, "Application of the leading-edge-suction analogy of vortex lift to the drag due to lift of sharp-edge delta wings," NASA Langley Research Center, Hampton, VA, TN D-4739, 1968.

[2] B. Singh, "Dynamics and aeroelasticity of hover capable flapping wings: Experiments and analysis," Ph.D. Dissertation, Depart. Aero. Eng., Univ. of Maryland at College Park, College Park, MD, 2006.

[3] B. Singh and I., Chopra, "Insect-based hover-capable flapping wings for micro air vehicles: Experiment and analysis," *AIAA Journal*, vol. 46, no. 9, pp. 2115- 2135, Sept., 2008.

[4] M. Benedict, M. Mattaboni, I. Chopra, and P. Masarati, "Aeroelastic analysis of a micro-air-vehicle-scale cycloidal rotor in hover," *AIAA Journal*, vol. 49, no. 11, pp. 2430-2443, Nov., 2011.

[5] R. Malhan, V. Lakshminarayan, J. Baeder, I. Chopra, P. Masarati, "CFD-CSD coupled aeroelastic analysis of flexible flapping wings for MAV applications: Methodology validation". *presented at 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, Honolulu, HI, 2012.

[6] R. Malhan, M. Benedict, I. Chopra, "Experimental investigation of a flapping wing concept in hover and forward flight for micro air vehicle applications," *presented at the 66th Annual American Helicopter Society Forum*, Phoenix, AZ, 2010.

[7] R. Malhan, "Investigation of aerodynamics of flapping wings for micro air vehicle applications," Ph.D. Dissertation, Depart. Aero. Eng., Univ. of Maryland at College Park, College Park, MD, 2013

[8] I. H. Tuncer, and M. F. Platzer, "Thrust generation due to airfoil flapping," *AIAA Journal*, vol. 34, no. 2, pp. 324–331, Feb, 1996.

[9] I. H. Tuncer, R. Walz, M. F. Platzer "A computational study of the dynamic stall of a flapping airfoil," *AIAA Journal*, pp. 2598–2519, June, 1998.

[10] K. Isogai, Y. Shinmoto, Y. Watanabe, "Effects of dynamic stall on propulsive efficiency and thrust of flapping airfoil," *AIAA Journal*, vol. 37, no. 10, pp. 1145–1151, Oct., 1999.

[11] A. Hart and L. Ukeiley, "Low Reynolds number unsteady aerodynamics over a pitching-plunging flat plate," *presented at 48th AIAA Aerospace Sciences Meeting*, Orlando, FL., 2010.

[12] W. Yuan, R. Lee, E. Hoogkamp, M. Khalid, "Numerical and experimental simulations of flapping wings," *International Journal of Micro Air Vehicles*, vol. 2, no. 3, pp. 181–209, Sept., 2010.

[13] S. Heathcote, Z. Wang, I. Gursul, "Effect of spanwise flexibility on flapping wing propulsion," *Journal of Fluids and Structures*, vol. 24, no. 2, pp. 183–199, Feb., 2008.

[14] S. K. Chimakurthi, J. Tang, R. Palacios, C. E. S. Cesnik, W. Shyy, "Computational aeroelasticity framework for analyzing flapping wing micro air vehicles," *AIAA Journal*, vol. 47, no. 8, pp. 1865-1878, Aug., 2009.

[15] P. Gopalakrishnan, D. K. Tafti, "Effect of wing flexibility on lift and thrust production in flapping flight," *Journal of Aircraft*, vol. 48, no. 5, pp. 865-877, May, 2010.

[16] D. B. Mayo, J. L. Lankford, M. Benedict, I. Chopra, "Experimental and computational analysis of rigid flapping wings for micro air vehicles," *Journal of Aircraft*, vol. 52, no. 4, pp. 1161-1178, July, 2015.

[17] D. B. Mayo, J. L. Lankford, M. Benedict, I. Chopra, "Aeroelastic analysis of avian-based flexible flapping wings for micro air vehicles," *Journal of the American Helicopter Society*, vol. 60, no. 3, pp. 1-18, July, 2015.

[18] D. Coleman, M. Benedict, V. Hrishikeshavan, I. Chopra, "Design, development and flight-testing of a robotic hummingbird," *presented at the American Helicopter Society International 71st Annual Forum*, Virginia Beach, VA, 2015.

[19] J. L. Lankford and I. Chopra, "A computational and experimental study of insect-based flexible flapping wing aerodynamics and structural deformation," *presented at the American Helicopter Society International 72nd Annual Forum*, West Palm Beach, FL, 2016.

[20] J. L. Lankford and I. Chopra, "A coupled CFD/CSD study of a flexible MAV-scale flapping wing in the hover flight condition," *presented at the American Helicopter Society International 73rd Annual Forum*, Fort Worth, TX, 2017.

Reduced-Order Model and Shaped Sensor for Flapping-Wing Control

Project: MCE-16-17-1.5

Mingjun Wei (PI), Haotian Gao, Wei Zhang
Mechanical and Nuclear Engineering Dept.
Kansas State University
Manhattan, KS 66506
mjwei@ksu.edu

Jayant Sirohi (Co-PI), Jason Tran
Aerospace Engineering and Engr. Mechanics Dept.
University of Texas at Austin
Austin, TX 78712
Jayant.sirohi@mail.utexas.edu

John Hrynuk, Christopher Kroninger
Vehicle Technology Directorate
U.S. Army Research Laboratory
Aberdeen Proving Ground, MD 21005

For the need of developing bio-inspired autonomous flapping-wing micro air vehicles (MAV), the current project explores the fundamental understanding and implementation of reduced-order models (ROM) for the real-time prediction of the coupled fluid-structure aerodynamics of flapping wings. The project also explores the possibility to design and build shaped sensors to provide efficient inputs for ROMs. In this report, the recent work by the research team is summarized to cover different but related subjects which all contribute to the overall goal of the project: the methodology and development of ROMs using POD-Galerkin projection on moving boundary problems (e.g. flapping wings), model adaptation for the generalization of ROMs from native parameters to new ones, the exploration and comparison of using DMD modes as base functions, ROM built and driven directly by experimental data, the design and test of shaped sensors on flapping wings. The conjunction and interaction between the areas are logically embedded in the work as well as the report.

***Index Terms*— reduced-order models, shaped sensors, flapping wings, data-driven models, POD-Galerkin projection, DMD, model adaptation**

I. Introduction

Development of practical bio-inspired flapping-wing MAVs requires optimization of the aerodynamic efficiency of flapping wings, accurate modeling of the vehicle flight dynamics and incorporation of a robust control system. Each of these tasks relies upon the calculation and even prediction of the loads generated by flapping wings across a range of operating conditions, such as flapping frequency, stroke kinematics and flight speed. Coupled fluid-structure modeling is required to capture the large deformation and nonlinear response of these highly flexible wings.

The calculation of the coupled aeroelastic response of a flapping wing becomes increasingly challenging as the flexibility of the wing and the flapping frequency increase. Accurate prediction of the unsteady aerodynamics and large structural deformations, especially at high flapping frequencies, requires refined coupled computational fluid dynamics and computational structural dynamics tools. These can yield accurate results but are computationally expensive and are not amenable to implementation for real-time control. While there have been many studies on the aerodynamics and structural dynamics of flapping-wing flight, stability and control aspects have only recently been explored [37,47].

The real-time knowledge of forces acting on the body of the MAV from the flapping wings can enable real-time optimization based predictive control techniques. These control techniques allow for full utilization of the flight envelope by maximizing the mission performance metrics within constraints, such as, mission requirements and physical capabilities of the vehicle. Predictive control requires real-time measurement of the forces acting on the vehicle, otherwise these forces would contribute to the uncertainty of the system, which ultimately translates to significant loss of performance [38, 39]. Hence real-time knowledge of the forces provides a means to effectively utilize predictive control to optimize flight performance. Therefore, there is a need for a control-oriented model that can predict the forces produced by the wings in real-time. In conjunction with onboard sensors that measure the state of vehicle and the wing deformation, the model can be used for effective stability and control.

The *research objective* of the current project explores reduced-order modeling to predict the coupled fluid-structure behavior of flapping wings. In parallel, shaped modal sensors are developed to real-time measure the deformed shape of the wing. In conjunction with the reduced-order model, these measurements are expected to yield the estimation and near-

future prediction of the instantaneous forces and moments produced by flapping wings [36,48].

The technical approaches involved in this project, as a collaborative work, cover quite a few areas involving a modified approach for POD-Galerkin projection for moving-boundary problems, DMD-based ROMs, modal analysis of experimental data, modal adaptation, and the design of shaped sensors for the measurement of specific modes, which are all summarized in this report.

II. Related Work

There have been *synergies with existing works and collaborations*. The methodology of global model reduction is tested on simulation database generated from a modified version of NMSU-FlexSI code [28, 29], which has been previously developed under Army High Performance Computing Research Center (AHPCRC). The model order reduction technique for moving boundaries is also studied for a related research under AHPCRC [12,22,50,51], though entirely different scenario (i.e. high-speed compressible flow) described by different governing equations is considered there. The understanding of flexible flapping wings benefits from the PI's participation in an AFOSR project, where both physics-based and adjoint-based approaches were taken to understand the impact of flexibility in flapping wings [29, 45]. Such understanding provided the rationality to include wing deformation in ROMs at the first place. The current research benefits from the communication and collaboration with Dr. Inderjit Chopra and Dr. James Baeder of University of Maryland and Mr. Pat Collins, Dr. Rajneesh Singh, Dr. Luis Bravo of ARL.

III. Technical Approaches and Results

The intrinsic complexity of the current project leads to a highly collaborative work with several key components. To make a clearer presentation of the work, in this session, four subjects are described separately: (A) POD-Galerkin projection for ROMs with moving boundary and ROMs with experimental data, (B) model adaptation, (C) model order reduction using DMD modes, and (D) design and test of shaped sensors.

A. POD-Galerkin projection for moving-boundary problems and to use experimental data

Part of the work in this session has been published and presented in conferences [23,32,49].

The traditional approach of POD-Galerkin projection for fixed fluid domain [1,2,3] is not capable to hand directly the moving-boundary problems (e.g. flapping wings). In an earlier study [23], we treated the combination of fluid and solid as one whole stationary domain, which was similar to the idea of immersed boundary technique used in numerical simulation [24-26]. The combined domain allows simple fixed meshes for discretization, while the effect of moving boundary/structure is represented as extra body-force terms added to the original Navier-Stokes equation on "solid" area. The modified Navier-Stokes equation for the combined domain is [27-29]:

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u}\cdot\nabla\mathbf{u} = -\nabla p + \frac{1}{\mathrm{Re}}\nabla^2\mathbf{u} + \mathbf{f}, \tag{1}$$

where $\mathbf{f}$ is the extra bodyforce term added in solid domain Ω_s to define the trajectory of solid boundary/structure.

With the modified fluid-solid equation (1), the classical approach for fixed fluid domain can be directly applied on the combined fluid-solid domain to obtain a dynamic equation,

$$\dot{a}_i = \sum_{j=0}^{N}(\frac{1}{\mathrm{Re}}l_{ij} - l'_{ij})a_j - \sum_{j=0}^{N}\sum_{k=0}^{N}(q_{ijk} - q'_{ijk})a_j a_k + c'_i, \qquad i = 1, ..., N, \tag{2}$$

which has some new parameters defined with support only in the solid domain:

$$l'_{ij} = \langle(\frac{1}{Re}\nabla^2\boldsymbol{\phi}_j + \frac{1}{\Delta t}\boldsymbol{\phi}_j), \boldsymbol{\phi}_i\rangle_{\Omega_s(t)}, \tag{3}$$

$$q'_{ijk} = \langle\nabla\cdot(\boldsymbol{\phi}_j\boldsymbol{\phi}_k), \boldsymbol{\phi}_i\rangle_{\Omega_s(t)}, \tag{4}$$

$$c'_i = \langle\frac{1}{\Delta t}\mathbf{V}^n, \boldsymbol{\phi}_i\rangle_{\Omega_s(t)}, \tag{5}$$

The resulted model (2) describes globally both fluids and solids. Different from the work by Liberge and Hamdouni [21], the current approach allows large and arbitrary movement of boundaries though only prescribed motion is considered at this moment.

Since the solid domain Ω_s is a time-dependent function, the integration in (3-5) needs to be computed at every time step. Though the computational process for simple integration is much less complex than the original high-fidelity model used in direct numerical simulation (DNS) and the solid domain is much smaller than the fluid domain, the increased computational cost by these terms are much larger than the ROM itself and overshadows the merit of using ROM. In our recent work [32], solid domain decomposition [21] is implemented to remove this extra cost.

Introducing a characteristic function,

$$\chi_s = \begin{cases} 1 & in \quad \Omega_s \\ 0 & otherwise \end{cases}, \tag{6}$$

like the fluid flow, the solid domain itself, marked now by χ, can also be decomposed to individual modes and the dimension is then reduced from a continuous space (or a space at the order of grid points in numerical simulation) to a low-order space with a truncated number of modes:

$$\chi_s(x,t) = \sum_{m=0}^{\infty} b_m(t)\psi_m(x) \approx \sum_{m=1}^{M} b_m(t)\psi_m(x), \tag{7}$$

where ψ is the solid modes computed by a similar POD approach but applied on domain function χ. The truncation of modes provides the good approximation of the same coefficients:

$$l'_{ij} \approx \sum_{m=0}^{M}\langle(\frac{1}{Re}\nabla^2\boldsymbol{\phi}_j + \frac{1}{\Delta t}\boldsymbol{\phi}_j), \psi_m\boldsymbol{\phi}_i\rangle b_m = \sum_{m=0}^{M} l^*_{ijm}b_m, \tag{8}$$

$$q'_{ijk} \approx \sum_{m=0}^{M}\langle\nabla\cdot(\boldsymbol{\phi}_j\boldsymbol{\phi}_k), \psi_m\boldsymbol{\phi}_i\rangle b_m = \sum_{m=0}^{M} q^*_{ijkm}b_m, \tag{9}$$

$$c'_i \approx \sum_{m=0}^{M}\langle\frac{1}{\Delta t}\mathbf{V}^n, \psi_m\boldsymbol{\phi}_i\rangle b_m = \sum_{m=0}^{M} c^*_{im}b_m, \tag{10}$$

where the new integrals, l^*_{ijm}, q^*_{ijkm}, and c^*_{im}, are not time-dependent and need to compute only once for each mode. The time-dependent variables are in b_m which are at much lower order (e.g. number of modes) than the original solid domain (e.g. number of mesh points in solid domain). Then, put these new terms in equation (2), get a new governing equation:

$$\dot{a}_i = \sum_{j=0}^{N}(\frac{1}{\mathrm{Re}}l_{ij} - \sum_{m=0}^{M} l^*_{ijm}b_m)a_j - \sum_{j=0}^{N}\sum_{k=0}^{N}(q_{ijk} - \sum_{m=0}^{M} q^*_{ijkm}b_m)a_j a_k + \sum_{m=0}^{M} c^*_{im}b_m. \quad (11)$$

To close the system, a level-set like equation is introduced to track the domain change [21] :

$$\dot{b}_n = -\sum_{n=0}^{N}\langle \boldsymbol{V}(x,t)\cdot\nabla\boldsymbol{\psi}_m, \psi_n\rangle b_m, \qquad n = 1, ..., N. \quad (12)$$

The modified POD-Galerkin projection has been applied to both two-dimensional and three-dimensional cases. For simplicity, this report only includes the three-dimensional case: a flow passing an oscillatory sphere.

The Reynolds number defined by the sphere diameter and incoming velocity is 300. The neutral position of the sphere center is at,

$$(x_0, y_0, z_0) = (0, 0, 0),$$

and the oscillation is defined by $y(t) = y_0 + A\sin(2\pi ft)$ with amplitude A=1.0 and frequency f=0.16. Figure 1 shows a typical DNS snapshot, where vortex structures are shedding behind the sphere.

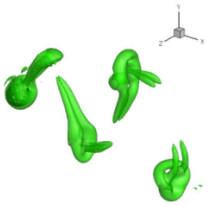
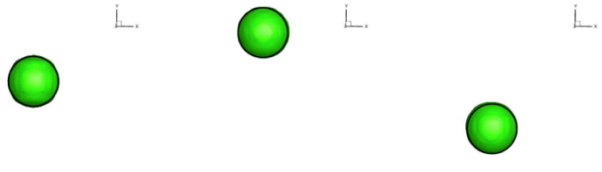


Fig. 1. A typical snapshot from direct numerical simulation: vortex structure of a flow passing an oscillatory sphere.

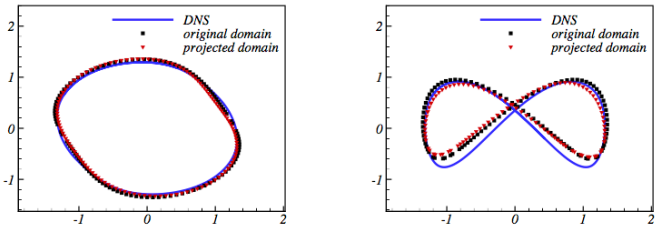
We first compare the accuracy of domain decomposition itself. Shown in figure 2, using 8 solid domain modes, we can accurately capture the shape and location of solid domain.



(a) (b) (c)

Fig. 2. Comparison of the exact shape and location of the solid (black circle) and the shape rebuilt form 8 solid modes (iso-surface in green) at different phases of one period: (a) t=0, (b) t=1/4T, (c) t=3/4T.

The dynamics of the global POD modes described in (11) has been captured successfully by the respective ROMs in a similar comparison of phase portraits of coefficients (figure 3).



(a) (b)

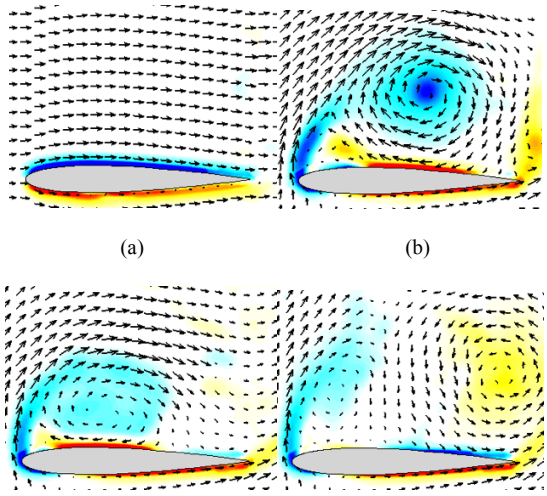
Fig. 3. Comparison of phase portraits of global POD coefficients computed by DNS, solid domain decomposition (i.e. projected domain), and continuous solid domain (i.e. original domain): (a) (a_1, a_2) and (b) (a_1, a_3)

Our recent study [32] also pushes the computational speed further by analyzing and removing some costly terms, which have negligible impact on the dynamics, from the ROMs. The time cost of different types of ROMs in comparison with the original DNS are shown in table I.

TABLE I. THE TIME COST OF FOUR PERIOD CALCULATION OF DIFFERENT METHODS FOR 8 MODES OSCILLATORY SPHERE CASE.

Method	Online-time	Offline-time
DNS	7 hours 15 mins	0
"easiest" method with Q terms	2 hours 8 mins	7 mins
Continuous method with Q terms	1-2 s	1 hour 54 mins
Decomposed method with Q terms	1-2 s	1 hour 1 mins
"easiest" method without Q terms	52 mins	7 mins
Continuous method without Q terms	1-2 s	22 mins
Decomposed method without Q terms	1-2 s	8.4 mins

It is appealing to avoid DNS entirely and build a ROM directly on experimental data. Figure 4 shows the experimental data of PIV snapshots generated in ARL's wind tunnel for the dynamic stall of a pitching-up airfoil provide, which have well-resolved velocity field (the same as numerical data). Body-fixed coordinate is used for convenience.



(a) (b)

(c) (d)

Fig. 4. The vorticity of the experimental pitching airfoil case. (a) t=0.21, (b) t=0.84, (c) t=1.47, (d) t=2.1.

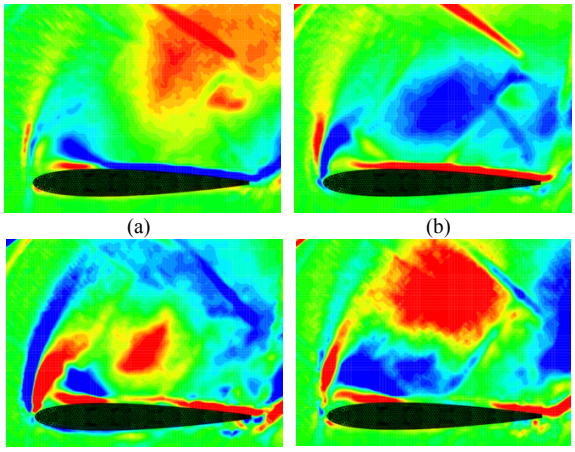


Fig. 5. The vorticity of the POD modes of experimental pitching airfoil: (a) mode 1; (b) mode2; (c) mode 3; (d) mode 4.

The first 4 POD modes, shown in figure 5, clearly capture the basic vortex structures behind the airfoil. In fact, the first 4 modes contain almost 90% of the total energy. And by using a ROM of only 8 modes, which have about 97% of the total energy, the lift and drag of the plate are successfully reproduced as shown in figure 6.

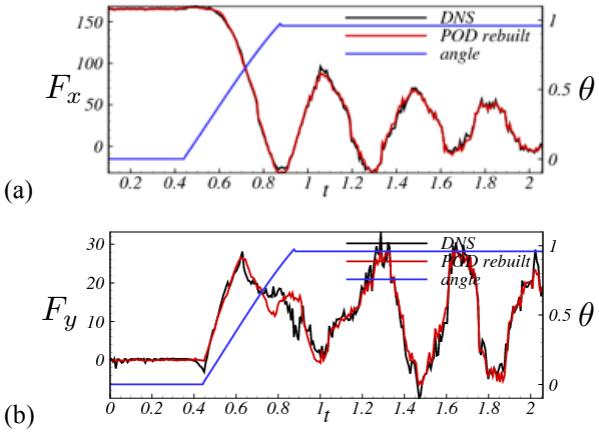


Fig. 6. The force results of the pitching airfoil, (a) drag, (b) lift.

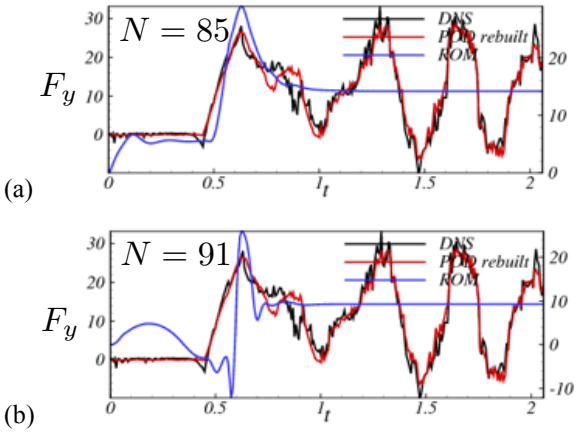


Fig. 7. The lift of the pitching airfoil when the ROM built by different number of snapshots, (a) N=85, (b) N=91.

The lift shown in figure 6(b) indicates that dynamic stall is happened at t = 0.63 where the angle of attack (AoA) is 26.2°. The prediction of the lift force and dynamic stall is plotted in figure 7. Using 85 snapshots (figure 7a), which cover the time from t = 0 to t= 0.59, the dynamic stall is accurately predicted with a good match of lift force; using 91 snapshots (figure 7b), which cover till the moment of dynamic stall, the stall angle is still predicted though the overall match of force value is not as accurate.

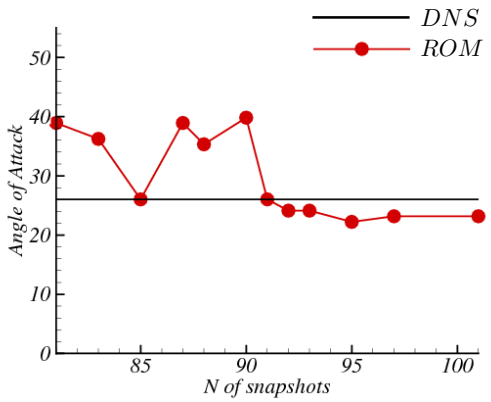


Fig. 8. The predicted AoA of dynamic stall changing with the number of snapshots which built the modes of the ROM.

This prediction continues to be accurate with more snapshots added (figure 8). However, when too many snapshots are used to build the ROM, the accuracy starts to drop and eventually (with snapshots covering more than t =1) the ROM fails to give a good prediction of the stall angle. The likely reason is that the dynamical behavior of the flow passing a pitching-up airfoil is "diluted" by the dynamics (i.e. snapshots) of a later stage when the pitching-up is all completed and the flow passing a fixed-airfoil (at high AoA) behaves differently and dominated by periodic vortex shedding.

B. Model adaptation

The ability to use a ROM with parameters at one set of "native/original" value to predict results under a different set of "new" parametric value is critical for the whole research on model order reduction. Often, a small deviation of parameters from their original/native values leads to inaccurate or even unstable results. The concept of model adaptation is to allow ROMs to be used for new situation with non-native parameters. Typical modal adaptation is to extend for different Reynolds numbers or Mach numbers [3,4,46], the extension to different motions is more challenging. One obvious reason is that, unlike Reynolds number or Mach number, the motion parameters do not explicitly appear in flow equations as well as the typical ROM built form a projection of the flow equations. The current formulation of using body force terms in flow equations to represent the solid motion makes possible to implement the motion parameters in flow equations. Introducing information from other states to modify and strengthen the ROM using one single set of data can potentially make the model more adaptable. In our study, two approaches using respectively the

idea or "shift mode" [3] and the idea of "combined model" [4] have been explored and modified to make ROMs more stable and accurate in model adaptation for different motions.

1. *Shift mode*

Different from the original approach [3] which defines a shift mode from the difference between a mean flow and a steady flow, our new approach defines the shift mode from the difference of two mean flows at different Reynolds numbers, since a steady solution does not exist for moving boundary problems.

The new shift mode was first applied to model the flow passing a two-dimensional fixed cylinder. Using the mean flows at Reynolds number 100 and 150, a shift mode is defined to extend the ROM from its native Reynolds number 100 to Reynolds number 150. It is shown that the modified ROM (with shift mode) and the original ROM (without shift mode) both work well at their native Reynolds number 100 in figure 9(a). However, when the ROMs are applied to Re = 150, the normal ROM without shift mode becomes unstable and the modified ROM is stabilized by the shift mode.

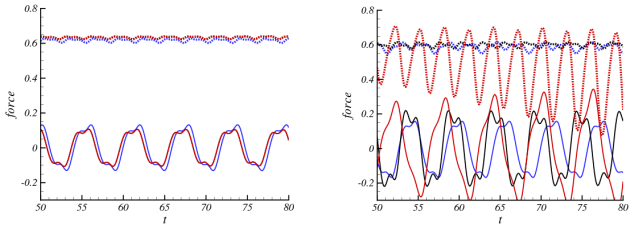


Fig. 9. Lift (solid line) and drag (dashed line) of the flow passing a cylinder: black lines are provided by the native normal ROMs at Re = 100 and 150 respectively for (a) and (b), red lines are provided by the normal ROM (generated at native Re= 100), and the blue lines are provided by the modified ROM with shift mode (generated at the native Re = 100, but with the mean flow at Re = 150): (a) Re =100; (b) Re = 150.

Then, the approach was applied to model a moving boundary problem: the flow passing a three-dimensional sphere where the sphere is oscillating normally along y-direction in a sinusoidal motion,

$$y = a\sin(2\pi f t).$$

The original ROM is based on the modes from the flow with amplitude a=1, and the shift mode was defined by the difference of mean flows of a=1 and a=0.5. For a quantitative comparison, an error function and a time-averaged error function are defined by

$$error(t) = \frac{\|\boldsymbol{u}(t) - \boldsymbol{u}_d(t)\|}{V_\Omega},$$

$$error_m = \frac{1}{T}\int_0^T \frac{\|\boldsymbol{u}(t) - \boldsymbol{u}_d(t)\|}{V_\Omega} dt,$$

where u(t) is the flow field predicted by ROMs, and $u_d(t)$ is the original DNS data. Figure 10 shows the error function defined to match the flow with an oscillation amplitude a=0.5 to compare the performance of ROMs in their prediction of the flow velocity. There are three ROMs used in this comparison: the native ROM using data at a=0.5 (as the benchmark), the direct use of the ROM at a=1.0 to predict the flow with a=0.5, and the ROM at a=1.0 but with a shift mode to predict the flow with a=0.5.

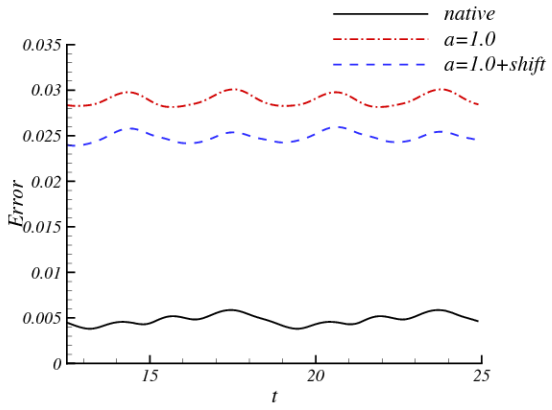


Fig. 10. The error changing with time for a native ROM at a=0.5, a normal "a=1" ROM, and an "a=1" ROM with a shift mode for their prediction of the flow with the oscillation amplitude a=0.5.

Unlike the earlier case for the adaptation of Reynolds number, the shift mode shows very limited help to improve a normal ROM in the adaptation to different amplitudes of the oscillation.

2. *Combined model*

The idea of combined model is to build the ROM based on the combined POD modes which are computed from all snapshots of different parameters. For the above problem of the oscillatory sphere, two sets of snapshots for the motion with a=0.5 and a=1 were included to compute the combined POD modes. Figure 11 shows the difference, as well as similarity, between the combined modes and the original POD modes at a=0.5 and a=1.

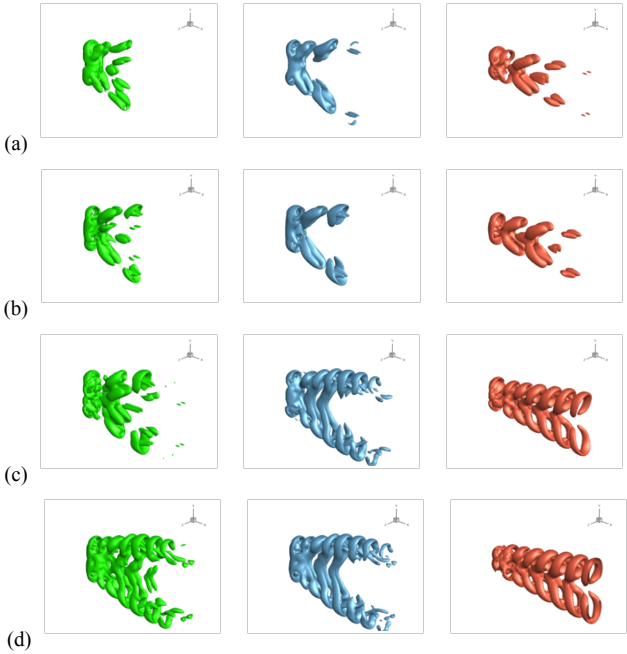


Fig. 11. The iso-contour for (left) combined modes, (center) a=1 modes and (right) a=0.5 modes. The first four modes are plotted in (a)-(d).

Based on the combined modes, the constructed combined ROM, using data of a=0.5 and a=1, was used to predict the flow for any oscillatory amplitude within the range [0.5, 1], and the prediction is compared to the prediction from native models at those amplitudes and the prediction of normal POD models using single set of data at either a=0.5 or a=1 (figure 12).

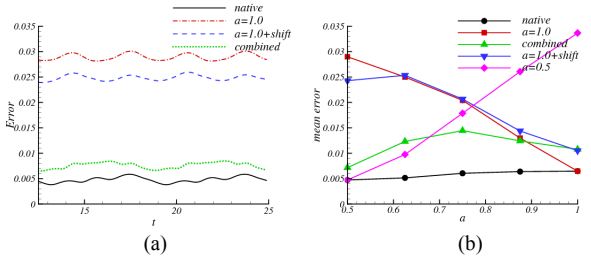


Fig. 12. (a) The error of different ROM to predict the flow of a=0.5, (b) the mean error of different ROM to predict the flow of other amplitudes a between 0.5 and 1.

The combined model shows much better prediction in its comparison to a normal single-parameter POD-Galerkin ROM and the ROM with a single shift mode, as shown in figure 12(a). A more thorough comparison is given in figure 12(b), where five different amplitudes were compared. The error of single-parameter ROMs, for either a=0.5 or a=1, increases quickly as the amplitude deviates further from their native parameters. The shift mode does not help much to reduce the derivation, even at the cost of less accuracy at the native parameter. The combined model, though not as accurate as native models, keeps the error low in its adaptation to any amplitudes between its "native" two amplitudes (a=0.5 and a=1). In practice, it is not possible to get native ROMs for all parameters, which require computation or experiments for all the parameters and defeat the purpose of ROM. The combined model, with only two data sets, shows an impressive capability in its prediction of other parameters for the model adaptation purpose to predict unknown states.

C. Model order reduction using DMD modes

Dynamic mode decomposition (DMD) has emerged as a new tool for the understanding of flow dynamics associated with frequencies [33,34,35]. The DMD modes computed by this process is considered here as an alternative of base functions for model order reduction. However, DMD modes are not orthogonal bases which are usually desired for the simplicity of Galerkin models. Therefore, the bi-orthogonal pair of DMD modes and adjoint DMD modes are both needed to build a simple DMD-Galerkin projection model. The introduction of adjoint DMD modes also provides an easy way to rank DMD modes for order reduction.

The approach is applied on a flow-passing-cylinder case in both transitional (figure 13) and periodic (figure 14) stages, where DMD and POD modes are compared. In both stages, DMD provides clear frequency information in the eigenvalues (figure 13a and 14a) and the time history of mode coefficients (figure 13b and 14b), which are not guaranteed by POD.

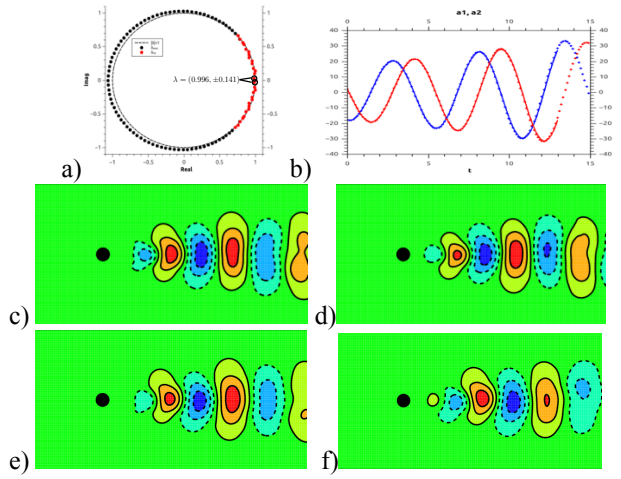


Fig. 13. Flow passing cylinder at the transitional stage: a) eigenvalues; b) time coefficient of the first DMD modes (real and imaginary parts in pair); c) real part of the first DMD mode; d) imaginary part of the first DMD mode; e) the first POD mode; f) the second POD mode

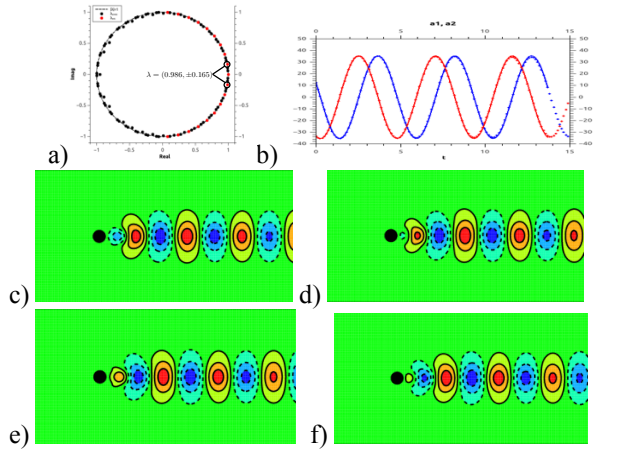


Fig. 14. Flow passing cylinder at the periodic stage: a) eigenvalues; b) time coefficient of the first DMD modes (real and imaginary parts in pair); c) real part of the first DMD mode; d) imaginary part of the first DMD mode; e) the first POD mode; f) the second POD mode

In figure 13 (a) and figure 14 (a), the most energetic modes are marked by red dots (19 for the periodic case and 29 for the transition case), and those are the ones used to build the ROM. Then, the temporal coefficients of the first mode are plotted in figure 13 (b) and 14 (b) for a comparison of the direct project from DNS (solid line) and the computation from ROM (dots). The real and imaginary components of the first DMD mode are plotted in pair (figure 13c, 13d, and figure 14c, 14d) and closely assemble the first pair of POD modes (figure 13e, 13f, and figure 14e, 14f) for both transitional and periodic cases.

For the periodic case, the DMD-Galerkin model performs similarly to the POD-Galerkin model; and for the transitional case, the DMD-Galerkin model carries sharper frequency features. On the other hand, the POD-Galerkin ROMs usually need less number of modes than the DMD-Galerkin ROMs to reach the same accuracy, and the POD-Galerkin ROMs are more robust in general.

D. Design and test of shaped sensors

Piezoelectric modal sensors were designed to measure the deflected shape of a flapping beam in real time. Based on a reduced-order model of the flapping beam, the measured deformation will be used to estimate the loads acting on the beam. This can be extended to a flapping wing, and will enable the real time, on-board measurement of wing loads to enable improved control algorithms on flapping wing micro aerial vehicles [36]. The real-time deformation measurement is enabled by shaped sensors bonded to the flapping beam. The shaped sensors perform modal filtering, and their output is combined to recreate the deformation. In this way, the deformation over the entire span of the beam is estimated from only a few sensor measurements. The shaped sensors are fabricated out of piezoelectric Polyvinylidene Flouride (PVDF) film [41, 42, 43, 44].

An experimental setup was constructed to impart translational and flapping motion to the beams and to measure the resulting deformation. Two shaped sensors, targeting the first and second natural modes of the beam, were bonded on the two sides of the beam. The deformation was also measured using Digital Image Correlation (DIC) [40] and the results were used to calibrate the piezoelectric sensors.

1) Theoretical basis

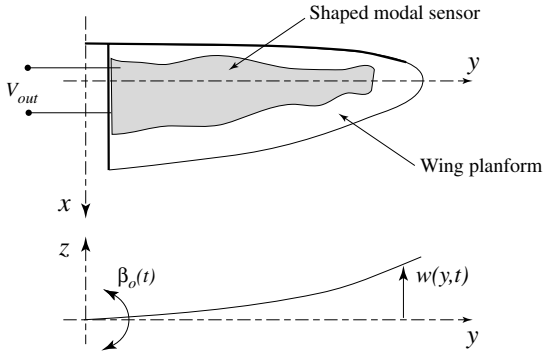


Fig. 15. Schematic of flapping wing with shaped piezoelectric sensor and associated coordinate system.

A schematic of a flapping wing with a shaped piezoelectric modal sensor is shown in figure 15. A flapping displacement $\beta(t)$ imparted to the wing root results in rigid flapping of the wing on which are superimposed elastic bending and elastic torsion. These elastic deformations are purely passive and occur in response to inertial and aerodynamic loads.

The transverse displacement $w(y,t)$ and the elastic twist angle $\theta_e(y,t)$ can be expressed in terms of a summation of shape functions as

$$w(y,t) = \sum_{i=1}^{N} \phi_i(y) q_i(t) = \boldsymbol{\phi q}$$

$$\theta_e(y,t) = \sum_{j=1}^{M} \eta_j(y) r_j(t) = \boldsymbol{\eta r}$$

where ϕ_i and η_φ are the assumed displacement shapes in bending and torsion respectively. The generalized coordinates in bending and torsion are q_i and r_j respectively. The equations of motion of the system, are given by

$$\begin{bmatrix} M_q & M_{qr} \\ M_{qr}^T & M_r \end{bmatrix} \begin{Bmatrix} \ddot{q} \\ \ddot{r} \end{Bmatrix} + \begin{bmatrix} K_q & 0 \\ 0 & K_r \end{bmatrix} \begin{Bmatrix} q \\ r \end{Bmatrix} = \begin{Bmatrix} Q_1 \\ Q_2 \end{Bmatrix}$$

where Q_1 and Q_2 are the generalized forces due to aerodynamics (functions of lift and pitching moment). Depending on the type of aerodynamic analysis used, the terms Q_1 and Q_2 will contain the degrees of freedom as well as their first and second derivatives. We propose to solve the inverse problem, i.e., given q and r, we will estimate the forcing. The shaped sensors will directly measure the modal coordinates q and r.

2) Non-rotating cantilevered beam

Consider a PVDF sensor bonded on the surface of a uniform, non-rotating cantilever beam that is undergoing pure bending displacement (replace the flapping wing in figure 9 with a cantilever beam). Assuming Euler-Bernoulli beam bending, and assuming uniaxial strain on the beam surface, the charge output $q(t)$ of the PVDF sensor is

$$q(t) = -d_{31} E_p \frac{t}{2} \int_0^L w''(y,t) b(y) dy$$

where d_{31} is the piezoelectric coefficient that relates electric charge to mechanical stress in the assumed uniaxial strain direction, E_p is the Young's modulus of the PVDF, t is the beam thickness, L is the sensor length (equal to the beam length), y is the lengthwise direction, and $b(y)$ is the sensor width. The bending displacement can be separated into a summation of mode shapes and modal coordinates. If the sensor is shaped (or alternatively, the electrodes are shaped) such that the sensor width $b(y)$ is equal to the second derivative of the sth mode shape, the charge from a sth modal sensor is

$$q_s(t) = C_s \sum_{r=1}^{\infty} \eta_r(t) \int_0^L \phi_r''(y) \phi_s''(y) dy$$

where C_s is a constant. Assuming that the PVDF sensor has negligible effect on the mass and stiffness of the beam, and making use of the orthogonality relationship between mode shapes, the charge output from the sth modal sensor is proportional to the sth modal coordinate.

3) Flapping wing

In the case of a wing undergoing flapping motion (simulated by a cantilever beam rotating about its root), the angular velocity results in a centrifugal stiffening, and the standard orthogonality condition for a uniform, non-rotating cantilever beam is no longer applicable. The equation of motion for a beam undergoing rotation about one of its ends, by a flap angle $\beta(t)$, can be derived as

$$\int_0^L EI \phi_r'' \phi_s'' dy + \frac{L^2 \dot{\beta}^2}{2} \int_0^L m_o \phi_r' \phi_s' \left(1 - \frac{y^2}{L^2}\right) dy = (\omega_r^2 + \dot{\beta}^2) \delta_{rs}$$

where ω_r is the natural frequency of the rth mode. For a beam with constant flexural rigidity, the orthogonality relation can be used to derive a relation between the modal coordinates and the charge developed by the PVDF sensor. Assuming N modal

coordinates are sufficient to represent the deformed shape to the desired accuracy, the charge developed can be written as

$$q_s(t) = C_s \sum_{r=1}^{N} \gamma_{sr} \eta_r(t)$$

where

$$\gamma_{sr} = (\omega_r^2 + \dot{\beta}^2)\delta_{rs} - \frac{L^2 \dot{\beta}^2}{2} \int_0^L m_o \phi_r' \phi_s' \left(1 - \frac{y^2}{L^2}\right) dy$$

This equation shows that in the case of a flapping wing, all the modes contribute to the charge developed by a sth modal sensor. However, using N modal sensors, it is still possible to extract the N modal coordinates at all times by solving the system of simultaneous equations.

4) Experimental setup

The piezoelectric sensors were fabricated from 28 μm thick PVDF film that is sputtered with 80 nm thick copper and nickel electrodes. The sensors were bonded to a carbon fiber beam of dimensions 101x26x0.4 mm. The beam with a first mode sensor bonded on its surface is shown in figure 16. A second mode sensor is bonded on the other surface of the beam. The shapes of these sensors are determined assuming an ideal cantilevered beam with uniform mass and stiffness distribution. In this way, the first and second modal coordinates can be measured simultaneously.



Fig. 16. First mode PVDF sensor on carbon fiber beam

The output of each sensor is conditioned by a circuit developed in-house, consisting of a charge amplifier followed by an inverting gain stage. The charge amplifier acts as a high pass filter and eliminates the effects of sensor and stray capacitance on the output. The circuit has a cutoff frequency of 0.3 Hz, which is sufficiently low to capture the first modal response of the beam. The cantilever beam was driven by an electrodynamic shaker. The beam was directly attached to the mounting point of the shaker, which provides a rigid body base motion to the cantilevered end of the beam. A geared servomotor was used to impart a flapping motion to the beam with a specified flap amplitude as a function of time.

5) Deformation measurement and modal identification

Digital image correlation (DIC) was used to measure the bending deformation of the cantilevered beam and flapping wing. In the DIC technique, a random speckle pattern is painted onto the specimen surface. Images are taken of the speckled specimen under no load (no motion) and under load (undergoing motion). The displacement of the speckle pattern is calculated by cross-correlating the unloaded and loaded images. The DIC yields a full-field measurement of bending deformation as a function of time, which is used to identify the modal parameters.

Modal parameters were identified with the Natural Excitation Technique (NExT) and the eigensystem relization algorithm (ERA). These techniques assume a linear, time-invariant structural model. The ERA is designed only for use with impulse excitation. Therefore, the NExT was first used to simulate a free response from a random input excitation; these free responses were then processed with the ERA to find the system matrix and output matrix of the state-space representation of the system. The mode shapes, natural frequencies, and damping ratios of the system were then calculated from these matrices.

Because the DIC measures full-field displacement, the calculated mode shapes are two-dimensional. These mode shapes are normalized such that the inner product of a mode shape with itself is equal to unity. Because the mode shapes are orthonormal, any modal coordinate q_r can be extracted by multiplying the measured displacement by $\phi_r(x, y)$ and integrating over the entire area.

The modal coordinates obtained from the DIC measurements were plotted against the outputs of the PVDF sensors. In this way, the PVDF sensor output voltages can be calibrated to the amplitude of the modal coordinates. The advantage of performing this calibration is that the piezoelectric constants and beam parameters do not have to be known exactly.

6) Results and discussion

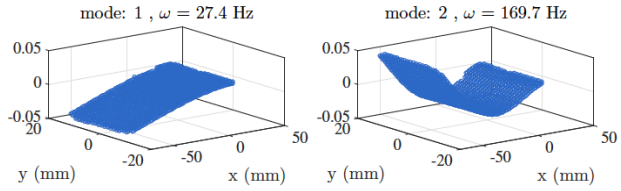


Fig. 17. Cantilever beam first and second mode shapes and natural frequencies, extracted from DIC measurements.

The cantilevered beam attached to the shaker was excited by a periodic random vibration signal. The signal was low pass filtered before input to the shaker so that only the first two modes of the beam were excited. The first and second mode shapes and natural frequencies of the beam, extracted from DIC measurements of bending displacement, are shown in figure 17. The extracted first and second natural frequencies (ω_1 = 27.4 Hz and ω_2 = 169.7 Hz) closely match the natural frequencies measured using a laser displacement sensor (ω_1 = 27.0 Hz and ω_2 = 169.0 Hz), thus validating the methodology for extraction of modal parameters.

7) Modal filtering

Figure 18 shows linear regression fits between the modal coordinates extracted from DIC (on the y-axis) and the output of the PVDF modal sensors (on the x-axis). It can be seen that there is high correlation between the first modal coordinate and the output of the first PVDF modal sensor, while there is very low correlation between the first modal coordinate and the output of the second PVDF modal sensor. Similarly, there is high correlation between the second modal coordinate and the output of the second PVDF modal sensor, while there is very low correlation between the second modal coordinate and the ouput

of the first PVDF modal sensor. This demonstrates that the modal sensors are succesfully filtering out their corresponding modal coordinates. The slopes of these linear fits can also be used as calibration factors between the output of the PVDF sensor and its corresponding modal coordinate.

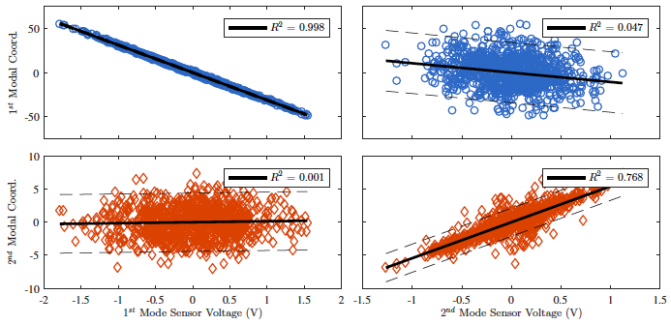


Fig. 18. Calibration of cantilever beam PVDF sensors to DIC modal coordinates from a periodic random vibration test.

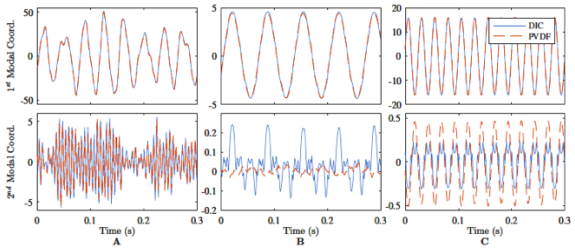


Fig. 19. DIC vs. PVDF measured modal coordinates for a cantilevered beam undergoing pure translation. A) Periodic random vibration B) 25 Hz pure sinusoid C) 40 Hz pure sinusoid.

Further tests were performed on the cantilever beam by exciting it with: (A) a periodic random signal, (B) a pure sinusoid at a frequency of and 25 Hz, and (C) a pure sinusoid at a frequency of 40 Hz. Figure 19 compares the modal coordinates extracted from DIC measurements to the modal coordinates measured by the PVDF sensors, i.e., the output of the PVDF sensors multiplied by their corresponding modal calibration factors. In general, the results show good comparison between DIC and PVDF measured modal coordinates, especially for the random vibration excitation. The second modal coordinate does not correlate as well as the first modal coordinate in the case of the pure sinusoidal excitation because the tested input frequencies are far from the second natural frequency and the resulting modal coordinates q_2 are small.

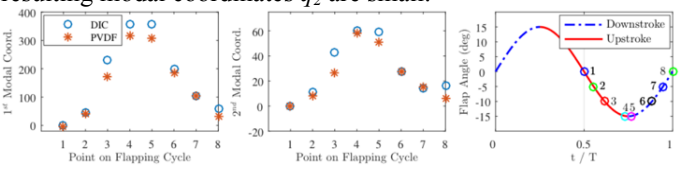


Fig. 20. DIC vs. PVDF measured modal coordinates for a cantilevered beam undergoing flapping motion.

Flapping wing experiments were then performed on the carbon fiber plate using mode shapes extracted from the random vibration testing. A flapping wing motion profile (12 Hz, 15° stroke amplitude) was used to compare DIC modal coordinates and piezoelectric modal coordinates. The resulting data is shown in figure 20. The trend of the modal coordinates measurements compares well for both modes. The points of largest error are around stroke reversal where deflection is greatest.

8) Modal sensor shape sensitivity analysis

It was noticed that there is some leakage between the modes as measured by the modal sensors. To identify the source of this leakage, the effect on modal sensing due to imperfections in construction was quantified by numerically calculating the output of sensors with perturbed shapes. The input deflected shape (w_i) is prescribed by choosing q_1 and q_2 values. The modal coordinates predicted by the perturbed sensors are then calculated, and the reconstructed deflected shape (w_r) is compared to the actual input shape. A measure of error is defined as

$$\%\ \text{error} = \left[\int_0^L (w_i - w_r)^2 dx \Big/ \int_0^L w_i^2 dx\right]^{1/2} * 100$$

The following perturbations, or construction errors, in the sensor shapes were analyzed:

1. Horizontal offset: The root of the sensor is offset from the cantilever root.
2. Vertical offset: The entire sensor is offset along the width of the beam. As a result, some regions of the sensor are no longer attached to the beam. The horizontal and vertical offsets simulate the effect of errors in bonding the sensor at the exact location.
3. Small width cut off: Width of the sensor that is smaller than a specified value is assumed to be zero. This simulates the effect of a finite thickness cutting tool.
4. Jagged edge: A random function generator is used to create a jagged edge. Parametrized by ratio between error from ideal shape and the beam width.
5. Linear distribution: This models errors in cutting out the shape of sensors by introducing a linear distribution of error with a maximum at the beam midpoint.
6. Array resolution: Models the method of creating a modal sensor by using discrete uniform pieces whose output signals are weighted to simulate a shaped sensor. Parametrized by array horizontal resolution divided by beam length, which defines the size of each piece in the array.

The sensitivity of a cantilevered beam modal sensor to these errors in construction is shown in figure 21. The ratio of the input modal coordinates was varied by increasing q_2 and keeping q_1 constant. The perturbations were normalized by a corresponding scale so that the relative significance of each construction error can be compared.

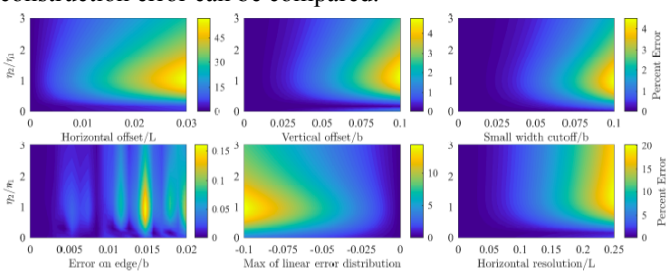


Fig. 21. Sensitivity of PVDF modal sensors to construction errors.

It is seen that the deformation sensing error due to a horizontal offset is the most significant, as a 3% perturbation in horizontal position can result in up to 50% sensing error. This is because the highest strain due to bending in a cantilevered beam occurs at the root. A vertical offset and small width cutoff model similar effects, so the significance of both is nearly identical. The modeling of a jagged edge produced nearly zero error in shape prediction. This is contrary to previous studies by other investigators who hypothesized that the reason for modal leakage in their sensors was due to the jaggedness of the sensor edge. The present numerical simulations imply that the modal leakage is most likely due to errors in horizontal placement of the sensor than errors in the actual shape of the sensor itself.

The sensitivity study of the piezoelectric array shaped sensor shows that it may be a promising concept for lower mode shape sensing. For a piezoelectric array composed of 10 pieces along the beam length, the error is reasonably small (less than 5%) for the values of q1 and q2 tested. As the resolution becomes more coarse, the modal filtering error increases.

The sensitivity plots also show that the maximum error is centered around $q_2 = q_1 = 1$. This is due to the mode 1 sensor being more sensitive to leakage at mode 2. The second mode sensor is more robust against construction effects in terms of filtering of the mode 1 coordinate. As q_2 / q_1 approaches 1, the mode 1 sensor is more heavily affected by q_2, which results in poor shape prediction. As q_2 / q_1 increases past 1, the output of the mode 2 shape sensor starts to dominate the signal, and the result is a more accurate prediction of the deformation. The construction errors have little effect on mode shape sensors sensing their own modal coordinate, but the effects are significant when other modal coordinates are present in the deflected shape.

9) Particle image velocimetry (PIV) measurements of a flapping beam

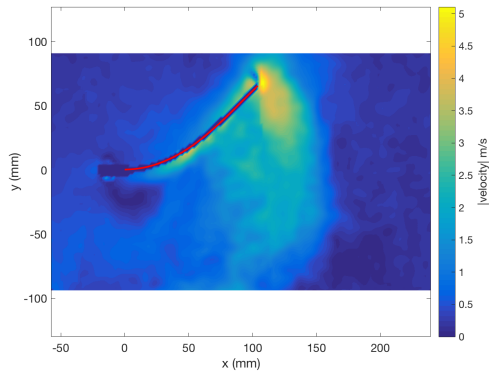


Fig. 22: Flow velocity magnitude field around flapping wing at stroke reversal. The beam edge is highlighted in red.

PIV measurement of a flapping beam was obtained to investigate the method of producing reduced order models from experimental data. The PIV experiment was performed on a 5-inch carbon fiber beam flapping in a 10 Hz sinusoidal motion. Phase-locked measurements were performed to obtain velocity field measurements at several discrete angles throughout the flapping cycle. Additionally, beam displacements were photogrammetrically obtained by beam edge illumination and loads were measured with a 6-axis load cell mounted at the root of the flapping beam. The dataset is useful it contains complete set of aeroelastic data for a flexible wing with a simple geometry.

10) Plate shape sensor prototype

Preliminary work has gone into the implementation of a shape sensor which may be able to reconstruct displacements of a plate which has a combined bending and torsional deformation. The plate shape sensor implementation is based on the same methods and tools which have been learned to implement the beam shape sensors. These include the use of digital image correlation for operational modal analysis and shape sensor calibration, the use of rectangular piezoelectric sensors, and the charge amplifier circuit. The main implementation challenge of the plate shape sensor deals with sensing of the first torsion dominated mode of the plate. This mode produces relatively little normal strains on the plate surface compared to shear strains, so the piezoelectric sensors must be oriented in a specific fashion to be sensitive to shear strains.

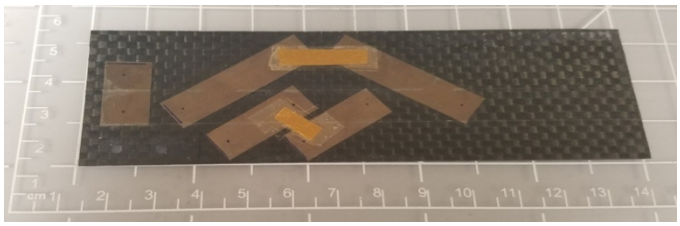


Fig. 23: Plate with rectangular piezoelectric sensors. Sensors are oriented to be sensitive to shear strains on the beam surface.

11) Beam shape sensor demonstration video

The experimental setup which was demonstrated at the MAST Capstone event has been documented in a video. The video displays real-time beam displacement reconstruction using rectangular piezoelectric sensors for flapping motion and deformation enforced by hand. The video is available online at https://www.youtube.com/watch?v=Y8tsPUnYCws

IV. Conclusion

There have been needs of reduced-order models for the purposes of real-time control and optimal design in both military and civilian applications. POD-Galerkin projection has been widely accepted as a mathematically rigorous approach to derive a low-order dynamic model for a complex dynamic system. When the models are implement for real-time control, appropriate sensors become critical to provide accurate inputs in an efficient manner. In particular, to develop a practical, autonomous flapping wing MAV, there is a need to develop reduced order models for real time calculation of the forces generated by the flapping wings, as well as sensors for the measurement of wing deformation that are compatible with the reduced-order models. This current work includes both the development of reduced-order models and appropriate sensors, which will also serve more general needs in related research areas in both Department of Defense and industries. The research provides numerical simulation data and experimental

data for the use of this project as well other projects of MAST and ARL. The comprehensive reduced-order models and the associated methodology are potentially available for technology transition for the use in the design, analysis and optimization of MAVs, and for the implementation in other systems requiring fast or real-time computation. A piezoelectric shaped sensor based on PVDF film, along with the associated instrumentation and signal-processing algorithms are developed in conjunction with the reduced-order model. This sensor system is potentially available for technical transition for deformation measurement when using a minimum of physical sensors is desirable.

REFERENCES

[1] Berkooz, G., Holmes, P., and Lumley, J. L., "The proper orthogonal decomposition in the analysis of turbulent flows," Ann. Rev. Fluid Mech., Vol. 25, 1993, pp. 539–575.

[2] Holmes, P., Lumley, J. L., and Berkooz, G., Turbulence, Coherent Structures, Dynamical Systems and Symmetry, Cambridge University Press, Cambridge, 1996.

[3] B. R. Noack, K. Afanasiev, and M. Morzynski, et al. "A hierarchy of low-dimensional models for the transient and post-transient cylinder wake," Journal of Fluid Mechanics, vol.497, No.1, 2003, pp.335-363

[4] Ma, X. and Karniadakis, G. E., "A low-dimensional model for simulating three-dimensional cylinder flow," Journal of Fluid Mechanics, Vol. 458, No. 1, 2002, pp. 181–190.

[5] Buffoni, M., Camarri, S., Iollo, A., Salvetti, M. V., et al., "Low-dimensional modelling of a confined three-dimensional wake flow," Journal of Fluiid Mechanics, Vol. 569, 2006, pp. 141–150.

[6] Lieu, T. and Farhat, C., "Adaptation of aeroelastic reduced-order models and application to an F-16 configuration," AIAA J., Vol. 45, No. 6, 2007, pp. 1244–1257.

[7] Siegel, S. G., Seidel, J., Fagley, C., Luchtenburg, D., Cohen, K., and McLaughlin, T., "Low-dimensional modelling of a transient cylinder wake using double proper orthogonal decomposition," Journal of Fluid Mechanics, Vol. 610, No. 1, 2008, pp. 1–42.

[8] Wei, M. and Rowley, C. W., "Low-dimensional models of a temporally evolving free shear layer," J. Fluid Mech., Vol. 618, 2009, pp. 113–134.

[9] Rowley, C. W. and Marsden, J. E., "Reconstruction Equations and the Karhunen-Lo`eve Expansion for Systems with Symmetry," Phys. D, Vol. 142, 2000, pp. 1–19.

[10] Rowley, C. W., Colonius, T., and Murray, R. M., "Model reduction for compressible flow using POD and Galerkin projection," Phys. D, Vol. 189, No. 1–2, Feb. 2004, pp. 115–129.

[11] Noack, B. R., Schlegel, M., Morzyn´ski, M., and Tadmor, G., "System reduction strategy for Galerkin models of fluid flows," International Journal for Numerical Methods in Fluids, Vol. 63, No. 2, 2010, pp. 231–248.

[12] Wei, M., Qawasmeh, B. R., Barone, M., van Bloemen Waanders, B. G., and Zhou, L., "Low-dimensional model of spatial shear layers," Physics of Fluids, Vol. 24, No. 1, 2012, pp. 014108.

[13] Cammilleri, A., Gu´eniat, F., Carlier, J., Pastur, L., Memin, E., Lusseyran, F., and Artana, G., "POD-spectral decompo- sition for fluid flow analysis and model reduction," Theoretical and Computational Fluid Dynamics, 2013, pp. 1–29.

[14] Tadmor, G., Lehmann, O., Noack, B. R., and Morzyn´ski, M., "Galerkin models enhancements for flow control," Reduced-Order Modelling for Flow Control, Springer, 2011, pp. 151–252.

[15] Noack, B. R., Tadmor, G., and Morzynski, M., "Actuation models and dissipative control in empirical Galerkin models of fluid flows," American Control Conference, 2004. Proceedings of the 2004, Vol. 6, IEEE, 2004, pp. 5722–5727.

[16] Tadmor, G. and Noack, B. R., "Dynamic estimation for reduced Galerkin models of fluid flows," American Control Conference, 2004. Proceedings of the 2004, Vol. 1, IEEE, 2004, pp. 746–751.

[17] Anttonen, J., King, P., and Beran, P., "POD-based reduced-order models with deforming grids," Mathematical and computer modelling, Vol. 38, No. 1, 2003, pp. 41–62.

[18] Anttonen, J., King, P., and Beran, P., "Applications of multi-POD to a pitching and plunging airfoil," Mathematical and computer modelling, Vol. 42, No. 3, 2005, pp. 245–259.

[19] Feng, Z. and Soulaimani, A., "Nonlinear aeroelasticity modeling using a reduced order model based on proper orthogonal decomposition," ASME Pressure Vessels and Piping Division Conference, No. 26006, 2007.

[20] Liberge, E., Benaouicha, M., and Hamdouni, A., "Proper orthogonal decomposition investigation in fluid structure interaction," European Journal of Computational Mechanics/Revue Europ´eenne de M´ecanique Num´erique, Vol. 16, No. 3-4, 2007, pp. 401–418.

[21] Liberge, E. and Hamdouni, A., "Reduced order modelling method via proper orthogonal decomposition (POD) for flow around an oscillating cylinder," Journal of Fluids and Structures, Vol. 26, No. 2, 2010, pp. 292–311.

[22] Wei, M. and Yang, T., "A global approach for reduced-order models of flapping flexible wings," AIAA paper 2010-5085, Chicago, IL, 2010.

[23] Gao, H. and Wei, M., "Global model reduction for flows with moving boundary," AIAA Paper 2014-0222, 2014.

[24] Peskin, C. S., "Flow patterns around heart valves: a numerical method," J. Comput. Phys., Vol. 10, No. 2, 1972, pp. 252–271.

[25] Peskin, C. S., "Numerical analysis of blood flow in the heart," J. Comput. Phys., Vol. 25, No. 3, 1977, pp. 220–252.

[26] Mohd-Yusof, J., "Combined immersed-boundary/B-spline methods for simulations of flow in complex geometries," Center for Turbulence Research, Annual Research Briefs – 1997, 1997.

[27] Zhao, H., Freund, J. B., and Moser, R. D., "A fixed-mesh method for incompressible flow-structure systems with finite solid deformations," J. Comput. Phys., Vol. 227, 2008, pp. 3114–3140.

[28] Yang, T., Wei, M., and Zhao, H., "Numerical study of flexible flapping wing propulsion", AIAA Journal, Vol. 48, No. 12, pp. 2909–2915, 2010.

[29] Xu, M., Wei, M., Yang, T., and Lee, Y. S., "An embedded boundary approach for the simulation of a flexible flapping wing at different density ratio", European Journal of Mechanics - B/Fluids Vol. 55, pp. 146–156, 2016.

[30] Brown, D. L., Cortez, R., and Minion, M. L., "Accurate projection methods for the incompressible Navier-Stokes equations," J. Comput. Phys., Vol. 168, 2001, pp. 464–499.

[31] Williamson, C., "Vortex dynamics in the cylinder wake," Annual review of fluid mechanics, Vol. 28, No. 1, 1996, pp. 477– 539.

[32] Gao, H. and Wei, M., "Domain decomposition in POD-Galerkin projection for flows with moving boundary," AIAA Paper 2016-1102, 2016.

[33] Mezic, I., "Spectral properties of dynamical systems, model reduction and decompositions," Nonlinear Dynamics, 41(1-3):309–325, 2005.

[34] Schmid, P. J., "Dynamic mode decomposition of numerical and experimental data," Journal of Fluid Mechanics, 656:5–28, 2010.

[35] Rowley, C. W., Mezic, I., Bagheri, S., Schlatter P., and Henningson D. S, "Spectral analysis of nonlinear flows," Journal of Fluid Mechanics, 641:115–127, 2009.

[36] Tran, J., Sirohi, J., Gao, H., and Wei, M., "Reduced order modeling of loads and deformation of a flexible flapping wing," AIAA Paper 2015-0177, 2015.

[37] Sun, M., "Insect flight dynamics: Stability and control," Reviews of Modern Physics, 86:615–646, May 2014. 1.

[38] Acikmese, B., Carson, J. M., and Bayard, D. S., "A robust model predictive control algorithm for incrementally-conic uncertain/nonlinear systems," International Journal of Robust and Nonlinear Control, 21:563–590, 2011.

[39] Carson, J. M., Acikmese, B., Murray, R. M., and MacMynowski, D. G., "A robust model predictive control algorithm augmented with a reactive safety mode," Automatica, 49:1251–1260, 2013.

[40] Sicard, J., and Sirohi, J., "Measurement of the deformation of an extremely flexible rotor blade using digital image correlation," Measurement Science and Technology, 24(6):065203, 2013

[41] Burke, S., and Hubbard, J. E., "Active vibration control of a simply supported beam using a spatially distributed actuator," IEEE Control Systems Magazine, 7(4):25–30, 1987.

[42] Gu, Y., Clark, R. L., Fuller, C. R., and Zander, A. C., "Experiments on active control of plate vibration using piezoelectric actuators and polyvinylidene fluoride (PVDF) modal sensorsm," Journal of Vibration and Acoustics, 116:303–308, 1994.

[43] Hubbard, J. E., "Spatial Filtering for the Control of Smart Structures: An Introduction," Springer-Verlag, Berlin, 2010.

[44] Lee, C. K., and Moon, F. C., "Laminated piezopolymer plates for torsion and bending sensors and actuators," Journal of the Acoustical Society of America, 85(6):2432–2439, 1989.

[45] Xu, M., and Wei, M., "Using adjoint-based optimization to study kinematics and deformation of flapping wings", Journal of Fluid Mechanics, Vol 799, pp. 56–99, 2016

[46] Qawasmeh, B., and Wei, M., "Low-dimensional models for compressible temporally developing shear layers", Journal of Fluid Mechanics, 731, pp. 364-393. 2013

[47] Hassanalian, M., Abdelkefi, A., Wei, M., and Ziaei-Rad, S., "A novel methodology for wing sizing of bio-inspired flapping wing micro air vehicles: theory and prototype", Acta Mechanica, 228(3), pp. 1097-1113, 2017

[48] Tran, J., Gao, H., Sirohi, J., and Wei, M., "Reduced-order methodology for prediction of loads generated by a flexible flapping wing", International Journal of Micro Air Vehicles, 2017

[49] Zhang, W., and Wei, M., "Model order reduction using DMD modes and adjoint DMD modes", AIAA paper 2017-3482, 2017

[50] Rezaian, E., and Wei, M., "Obtaining a stable Galerkin ROM in presense of shock-vortex interactions", AIAA paper 2017-1008, 2017

[51] Tabandeh, M., Wei, M., and Collins. J. P., "On the symmetrization in POD-Galerkin model for linearized compressible flows", AIAA paper 2016-1106

Bio-inspired Principles of Extended Terrain Mobility

MAST MCE-17-2.2

Robert J. Full, Kaushik Jayaram, Chen Li, Shilpa Naik, Nathaniel Hunt, Crystal Lee
Dept. of Integrative Biology,
University of California, Berkeley
rjfull@berkeley.edu

***Abstract*—Small animals show exceptional ability to traverse multi-component 3D terrain that far exceeds any robot yet built. Here, we continue our efforts at quantifying the mechanisms by which animals locomote in confined spaces, negotiate grass-like beams, and move along branches/rods/wires at high speeds. Cockroaches can ingress and egress into and out of almost any structure by using their body shape-changing ability. Previously, we discovered that cockroaches could traverse horizontal crevices smaller than a quarter of their body height in less than a second. Once inside vertically confined spaces, cockroaches were able to still locomote rapidly at 60 centimeters per second using an as yet unquantified mode of locomotion—body-friction legged crawling. Here, we provide the mathematical model that allows prediction of the performance of confined space locomotion between two horizontally stacked plates. Defining two locomotor regimes – stick, where foot does not slip on the ground, and slip where feet slip with respect to the ground surface permitted the prediction of velocity and stride length as a function of both ground and ceiling friction. We showed that cockroaches likely use vision when negotiating grass-like beams by recording their behavior when sight was blocked. Previously, we found that the American cockroach transitioned from an upright to an inverted running gait as it ran along rods of smaller diameter, remarkably maintaining high speed even during inverted locomotion. Here, we showed that despite only small changes in kinematics, cockroaches did employ their claws more during inverted running. We have begun to design a robot that could crawl along pipes, wires or conduits of damaged buildings. Lastly, we began a characterization of animal body morphology and natural terrain by using photogrammetry. Photogrammetry as well as other images based approaches combined with material property characterizations may someday lead to a terrain cloud database that could guide robot locomotion in the real environment.**

***Index Terms*—bio-inspiration, multi-component terrain, terradynamics, confined space locomotion, rod locomotion, photogrammetry, legged robots**

I. Introduction

A. Background and Technical Challenges

Natural terrain is complex. Not only does the ground surface often have challenging topology (different heights, slopes) and rheology (friction, stiffness, damping, viscosity), but it also has a variety of natural 3D components such as grass, stems, branches, trees, cracks and crevices in bark and leaf litter (Fig. 1) that go beyond rough terrain at the spatial scales of MAST Platforms (Fig. 2).



Fig. 1. Multi-component, 3D terrain common in nature which MAST scale robots will encounter (a) grass, (b) tree branches, (c) leaf litter, and (d) crevices in tree trunks.

Unfortunately, these natural multicomponent 3D terrains, as well as their human-made analogs that include walls, conduits, struts, pipes, vents and wires, can prevent current MAST Ground Platforms from executing the desired scenarios. Fortunately, animals negotiate these extended terrains often with ease and at high tempos. We have discovered fundamental biological principles of extended terrain mobility to inspire the design of the next generation of MAST Ground Platforms, so as to further develop their unique capabilities at this size scale.

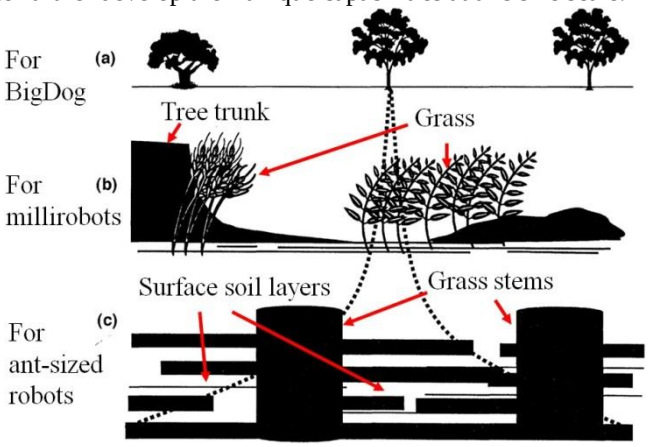


Fig. 2. Views of a natural terrain from the vantage point of (a) BigDog, (b) a current MAST scale robot, and (c) a future ant-sized robot.

We build on the success of MCE-2.2 Bio-inspired Foot and Appendage Design where we discovered the function of every structure (claws, sticky pads, ankles, spines and scales) on insect and lizard feet/legs/tails permitting running and climbing on varying surfaces. These biological principles were derived through close collaboration with MCE-2.1 Bio-inspired Dynamic Modeling and Simulation with Parameters for Ground Contact Model that has led to the first attempt at a Theory of Terradynamics [1] and translation to MAST Ground Platforms. Likewise, our discoveries of appendage function that include tails have led to a new field termed Inertial Assisted Control [2]. These discoveries are being translated to MAST robots in MCE-2.4 Bio-inspired Crawling, Running, Climbing Robots resulting in unprecedented terrain negotiation. The transformation from a structure-based focus to a capability-centered effort has allowed the realization of our vision for the biological inspiration of MAST Platform evolution.

B. State of the Art

At present, there exists no field of locomotor biology studying extended terrain mobility. There has been significant progress determining the biomechanics of how legged animals locomote on rough, granular, compliant and low contact probability terrain stemming from our initial studies [3-6] and others [7-8]. However, here our effort is among the first to systematically examine the capability of legged animals negotiating 3D multi-component terrains. We studied the mobility of small legged animals such as insects negotiating extended terrain in three scenarios: (1) locomotion over, through and along multi-component, 3D terrain analogous to grass (Fig. 3a), (2) confined space locomotion and ingress and egress through reduced spaces (e.g. cracks and crevices in leaf litter and tree trunks, Fig. 3b), and (3) vertical/horizontal branches (Fig. 3c,d).

C. Research Objectives

At least two major challenges must be met to quantify and model extended terrain mobility. First is the difficulty in creating models of multi-component, 3D terrain in the laboratory for systematic experimentation. New parameterization is required for components such as rows of vertically compliant beams (e.g. grass), horizontal beams (e.g. branches), confined spaces and narrow gaps (e.g. litter and trunks). Modeling these terrains will demand materials testing to quantify components and create **general extended terrain templates**. The second challenge is the ability to quantify the large variety of locomotor behaviors that small animals display (Fig. 3) so we can extract general biological principles that can be translated to the robots. First, we needed to explore and categorize behavioral strategies relevant to translation to robots. From these categories, we created a **generalized decision hierarchy for extended terrain** from which the most effective strategies are selected (e.g. plow through passively, climb over, run around, change locomotor strategy, or avoid route with high energy cost). To understand the mechanisms used, we needed to conduct experiments on the animals to determine which morphological features permit the novel capability and provide the information used in the decision. This required manipulating feet, legs, bodies, and sensors, as well as characterizing their material properties.

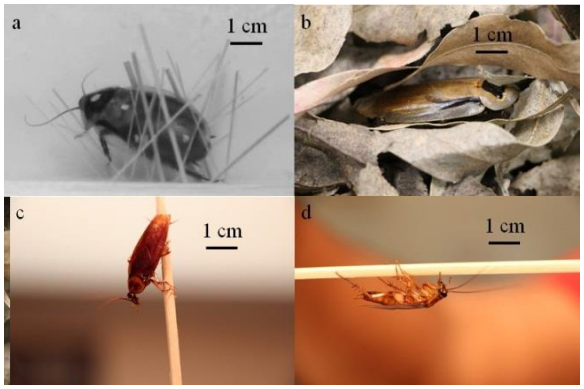


Fig. 3. Animals use a variety of locomotor modes in multi-component, 3D terrain. Cockroach maneuvers through a model "grass" field (a), wedges in layers of leaf litter b), climbs upside-down (c) and inverted (d) along a thin pole.

We hypothesized that the exceptional, and seemingly unstoppable, locomotor performance of animals were attained by: (1) using existing appendages in novel ways to sustain current strategy, (2) using alternate mechanisms that rely on interaction of the body and appendages with environment to adapt current strategy or generate new behavior, (3) physically conforming to the spatial environmental restrictions, and (4) modifying the environment, itself, to enable optimal use of existing locomotion strategies. We performed experiments in conditions that simulate extended terrain situations to elucidate novel, high-tempo locomotor behavior in animals.

II. Related Work

A. Synergy with Existing Works and External Collaborations

We continue to have very productive collaborations with the Goldman Lab at Georgia Tech. and the Fearing Lab at Berkeley. Students of the Full and Fearing Lab have continued collaborating on modifications of the OctoRoACH MAST platform. Inspired by our extended terrain mobility discovery in animals, we furthered collaboration with the Fearing Lab to develop a smooth shell for the VelociRoACH robot allowing extended terrain mobility.

B. Technology Transitions Summary

Jayaram and Haldane of the Fearing Lab at UC Berkeley have collaborated on the modification of the VelociRoACH hexapedal robot to run in confined spaces. Jayaram worked with Hunt to explore a robot that could move like a cockroach on a rod. Naik worked on using photogrammetry to obtain accurate 3D scans of natural terrain and animals.

III. Technical Approach

A. Locomotion in Confined Spaces

Small-sized animals and MAST robots have the opportunity to use solutions not available to larger platforms.

For example, a small insect or robots may have the capability to locomote in confined space. We showed that cockroaches can still locomote in vertically compressed spaces less than ¼ the animal's standing height [9] (Fig. 4).

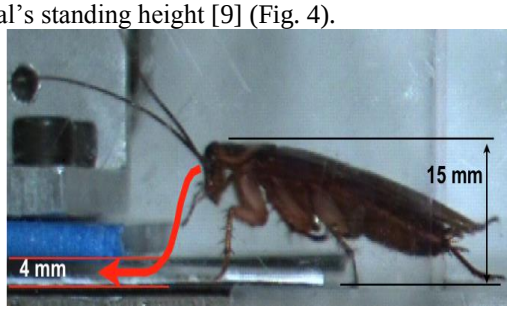


Fig. 4. Standing height relative to the gap size that still permits locomotion.

Cockroaches maintain their speed using a tripod gait and a constant duty factor with little slipping, despite being vertically compressed and locomoting with a far more sprawled posture. We varied vertical (Fig. 5) and lateral space, the space's shape and its frictional properties corresponding to pipes, conduits and vents. We measured metrics that included success rates, speed, gait, stride frequency, stride effectiveness (slipping) and tortuosity (straightness of the path). We also systematically changed surface properties (friction, texture).



Fig. 5. Experimental setup to confine vertical locomotion. Height and track surface properties (friction, texture) are easily modifiable and force sensors can be added.

From our discoveries, we were able to design a robot [9], CRAM (Crawling Robot with Articulated Microstructure) that realized this unique capability of confined space locomotion. Here, we effort to make a quantitative model that can predict performance of what we term frictional legged crawling.

B. Locomotion in Multi-component, 3D Terrain

We choose to study two types of multi-component, 3D terrain: (1) compliant vertical beams (grass), and (2) stiff poles/rods/wires (stems and branches).

We performed field surveys to determine the extent of natural variation in grass, branches and leaf litter to make our parameter variation most relevant. We varied terrain parameters (i.e. height, width, diameter, curvature, spacing, angle, stiffness, and friction), and observed from videos the success/failure rates of animals when moving through terrain in each condition and identified different locomotion strategies. We obtained high-speed, 3D kinematic data to characterize in detail the most successful locomotor modes for each condition. By this systematic approach, we generated specific hypotheses for effective locomotion through and along each terrain.

1) Vertical flexible beam obstacle negotiation.

To test our hypotheses concerning vertical beam negotiation, we modified the morphology and mechanics of body and sensory feedback of the animal, and tested if these changes enabled or disabled capability, and improved or decreased performance. The modification to body included adding an artificial shell to change the body shape and surface friction. Our experiments showed a novel sideways rolling behavior when cockroaches encounter a grass field (Fig. 6 a,b). Based on this observation, we found that the animal's slightly curved, oval-shaped, streamlined, smooth body aided negotiation [10].

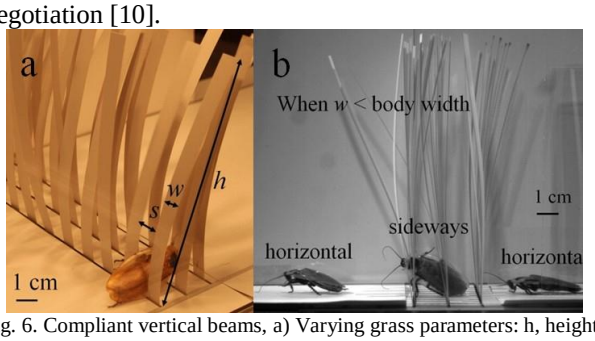


Fig. 6. Compliant vertical beams, a) Varying grass parameters: h, height; w, width, and s, lateral spacing. b) Lateral view of the sideways running rolling behavior when lateral spacing was smaller than body width.

We modified the shape of the animal's shell by adding an extra cornered, flat, square frame with higher friction, and measured the decrease in the animal's performance or failure to maneuver through grass-like beams. We conducted comparable experiments on the bodies of robots in collaboration with the Fearing Lab. Here, we began modification to sensors that included removing antennae and blocking eyes.

2) Stem/Branch/Pipe/Wire Rod Running

For locomotion on branches/stems/rods, we conducted a field survey to determine natural variation. We then created model poles of different diameter, orientation, and surface properties to model branches (Fig. 7).

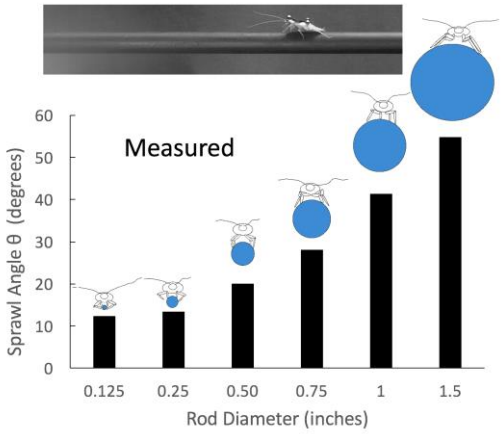


Fig. 7. Horizontal beam locomotion. (Top) Cockroach running along horizontal beam with markers for 3D kinematics. (Bottom) Sprawl angle of cockroach legs as a function of the rod diameters used.

Our initial observations led us to hypothesize that with decreasing pole diameter, animals will grasp the poles closer to the body center of mass to reduce destabilizing moments acting on the body. We previously discovered [11] that cockroaches can run equally fast upside down on the smallest stem-like rods and right side up on the largest branch-like rods. We hypothesized that this surprising inverted branch running capability may rely on a combination of passive stability from pendulum-like dynamics to maintain body angle and exploitation of foot morphology with context specific leg coordination to maintain attachment to the rod. Here, we chased cockroaches along the smallest diameter branches and examined the role of foot function during inverted running using high-speed kinematics.

C. Characterizing Multi-component, 3D Terrain Using Photogrammetry

To understand the challenge presented by real-world multi-component terrain to animals, and to visualize the interaction of their geometries, we obtained accurate 3D models of natural terrains and the animals. We found an inexpensive, mobile method to obtain 3D models from photographic images - photogrammetry. Using this technique, we took photographs of an object from different angles and then stitched them together to obtain the surface mesh of the object. This can be done by matching common features between successive images. A scale was placed next to the model for calibration of the mesh generated. For obtaining 3D models of the cockroaches, we mounted a freshly deceased animal on a stick and took photographs of it from different angles. For natural terrains, we took photographs of terrains outdoors from many angles in natural lighting.

Cloud based software (e.g. *Memento* Autodesk) allowed us to upload our images and the download the mesh generated from those images. Reconstructed 3D model's file size depended on the resolution of the input images and we have the capability (*Remake* Autodesk) to process even large files (10s of GB, 2+ billion polygon meshes). This software also allowed us to take sufficiently accurate measurements from the models by calibrating a fixed distance. The resolution of the meshes obtained was primarily dependent on the resolution of the images allowing us to obtain meshes for even our small-sized samples. Our experience revealed that we need approximately 100-150 high resolution images of the object to get a sufficient mesh for further analysis. To get high resolution images of our small sample, we used macro lenses at small apertures and high shutter speeds. The camera was fixed on a tripod, while the object was rotated on a turntable. The background of the object was kept plain white to limit distortion.

IV. RESULTS

A. Confined Space Locomotion

We continued our investigation of the mechanics of confined space locomotion using a physics-based single leg – compliant body model (Supporting Information [9]). For ease of analysis, we separated locomotion into two regimes – (1) Stick – where foot does not slip on the ground, and (2) Slip – where feet slip with respect to the ground surface. We also developed analytical relationships for three characteristic leg angles – slip, maximum, and minimum that determine these regimes.

During the stick regime, the body displacement is equal horizontal component of leg excursion. To determine mean body velocity during this phase, we developed an actuator model for the leg. Let us consider a simple, reciprocating leg actuator with a linearly decreasing force-velocity relationship limited in both force (F_{max}) and velocity (v_{max}) driven by a controller that attempts to output the maximum possible propulsive force during slip phase. Further, we assume that the leg forces are sufficient to overcome drag (at least at some leg orientation), but insufficient to compress the body (thus unable to reduce normal forces due to confinement).

$$\frac{F_L}{F_{max}} + \frac{v_L}{v_{max}} = 1 \tag{1}$$

Balancing forces on the body in the horizontal direction during stick mode,

$$\begin{gathered} F_L cos(\theta) = (\mu_c + \mu_g) N_g + \mu_c F_L sin(\theta) \\ \frac{F_L}{F_{max}} = \frac{C_h(\mu_c + \mu_g)}{cos(\theta) - \mu_c sin(\theta)} \\ \frac{v_L}{v_{max}} = 1 - \frac{C_h(\mu_c + \mu_g)}{cos(\theta) - \mu_c sin(\theta)} \end{gathered} \tag{2}$$

Further, the body velocity (v_b) and leg velocity (v_L) can be related using the geometric constraint below.

$$\begin{gathered} x^2 + h^2 = l^2 \\ \text{Differentiating, } x\dot{x} = l\dot{l} \rightarrow \dot{x} = \left(\frac{l}{x}\right)\dot{l} \\ v_b = v_L sec(\theta) \end{gathered} \tag{3}$$

Combining Eqn. 14 and Eqn. 15, we obtain the velocity of the body under stick (v_{stick}) as

$$\frac{v_{stick}}{v_{max}} = \left(1 - \frac{C_h(\mu_c + \mu_g)}{cos(\theta) - \mu_c sin(\theta)}\right) sec(\theta) \tag{4}$$

The mean velocity ($\tilde{v}_{stick}$) during stick mode is obtained by the following equation and is computed numerically.

$$\tilde{v}_{stick} = v_{max} \frac{\int_{\theta_{slip}}^{\theta_{max}} \left(1 - \frac{C_h(\mu_c + \mu_g)}{cos(\theta) - \mu_c sin(\theta)}\right) sec(\theta)\, d\theta}{\int_{\theta_{slip}}^{\theta_{max}} d\theta} \tag{5}$$

Additionally, the stick duration (t_{stick}) is

$$t_{stick} = \frac{x_{stick}}{\tilde{v}_{stick}} \quad (6)$$

The velocity at end of stick mode (or beginning of slip mode, v_0) is obtained as

$$v_0 = v_{max}\left(1 - \frac{C_h(\mu_c + \mu_g)}{cos(\theta_{slip}) - \mu_c sin(\theta_{slip})}\right) sec(\theta_{slip}) \quad (7)$$

Assuming that the leg slides back at the fastest speed it can once the foot starts slipping, we can estimate the net slip duration (τ_{slip}).

$$\tau_{slip} = \frac{h(cot(\theta_{min}) - cot(\theta_{slip}))}{2v_{max}\frac{\int_{\theta_{min}}^{\theta_{slip}} sec(\theta)\, d\theta}{\int_{\theta_{min}}^{\theta_{slip}} d\theta}} \quad (8)$$

We can compute the instantaneous and mean acceleration (a_{slip} and $\tilde{a}_{slip}$ respectively) on the body during the slip mode as

$$ma_{slip} = (\mu_f - \mu_c)F_L sin(\theta) - (\mu_c + \mu_g)N_g$$
$$\tilde{a}_{slip} = \frac{\int_{\theta_{min}}^{\theta_{slip}} \left((\mu_f - \mu_c)F_L sin(\theta) - (\mu_c + \mu_g)N_g\right) d\theta}{\int_{\theta_{min}}^{\theta_{slip}} d\theta} \quad (9)$$

Therefore, the instantaneous and mean body velocities (v_{slip} and $\tilde{v}_{slip}$ respectively) are computed as below.

$$v_{slip} = v_0 - \tilde{a}_{slip} t, \qquad t \le \tau_{slip}$$
$$t_{slip} = \frac{v_0}{\tilde{a}_{slip}}, \qquad t_{slip} \le \tau_{slip} \quad (10)$$
$$\tilde{v}_{slip} = \frac{v_0 - \tilde{a}_{slip} t_{slip}}{2}$$

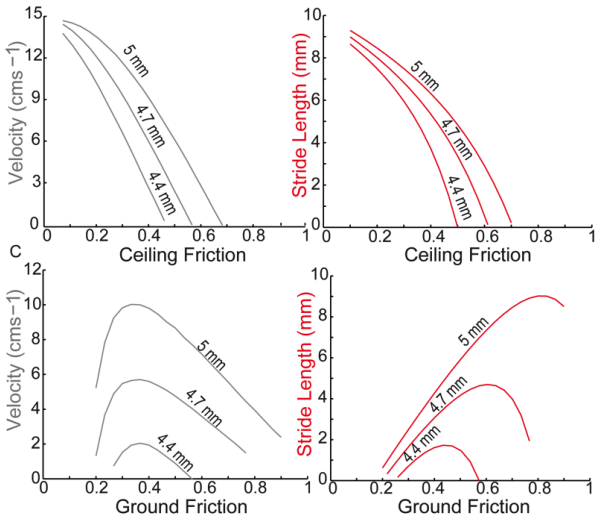


Fig. 8. Estimated velocity and stride length using single leg compliant body model for confined space locomotion.

Using the aforementioned equations, we calculated stride lengths and mean stride velocities during the stick and slip regimes, while varying ceiling and ground friction. These numerically estimated values are in close agreement with experimental observations (Fig. 8).

B. Locomotion in Multi-component, 3D Terrain

1) Vertical flexible beam obstacle negotiation

To gain initial insights into the role of sensors in negotiating multi-component 3D terrain, we began preliminary experiments to test the role of vision during cockroaches going through beam obstacles.

In our systematic study [10], we noticed that the cockroaches' obstacle traversal performance is still far superior to that of the robot, even when the robot is equipped with the cockroach-inspired rounded shell.

We hypothesized that sensory feedback may be important for the animal to initiate body roll or to use roll maneuvers that are the most effective. To begin to test this hypothesis, we explored the role of vision. We performed preliminary experiments using paint (whiteout) to cover the cockroach's eyes, including compound eyes for forming images and single eyes for sensing brightness. We observed that, as compared to an intact cockroach, which has a high probability of using roll maneuver, a blind cockroach more frequently explored laterally, but still could traverse the beam obstacles, although more slowly (Fig. 9).

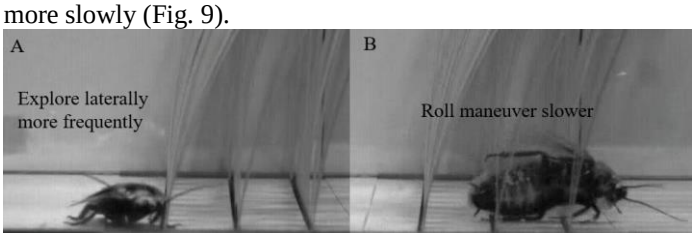


Fig. 9. Preliminary observations of a cockroach negotiating grass-like beam obstacles with vision removed.

We also performed preliminary experiments ablating the animal's major mechanosensing ability, including ablating one antenna and ablating both antennae. We observed that, as compared to an intact cockroach, which has a high probability of using roll maneuver, a cockroach missing one antenna more frequently explored laterally, and a cockroach missing both antennae more frequently climbed against and fell and flipped over. However, they still could traverse the beam obstacles, although more slowly (Fig. 10).

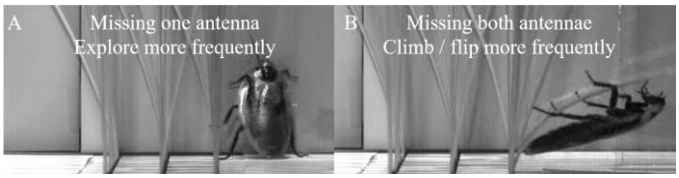


Fig. 10. Preliminary observations of a cockroach negotiating grass-like beam obstacles with antenna(e) ablated.

We do note that this observation is only preliminary, as we have yet to obtain a large enough sample size by testing multiple animals, analyze the video data to precisely measure traversal performance, and locomotor mode probabilities, and run statistical tests, to confirm this hypothesis. If our hypothesis is supported, it would provide inspiration for sensors to further improve our robot's performance to traverse densely cluttered obstacles, such as grass-like beams.

2) Stem/Branch/Pipe/Wire Rod Running

a) Foot use and leg coordination during inversion

We tested our hypothesis that foot use during inverted running played a key role in performance by using close-range high-speed video recordings. We observed that the multi-structured tarsus made substrate contact with both the claw and euplantulae or friction pads (Fig. 11a). Preliminary data supports the finding that these animals make use of both structures for running beneath large diameter rods, but when inverted running on smaller diameter rods they rely more heavily on claws rather than their euplantulae (Fig. 11b).

We have begun to track multi-camera high-speed videos to compute 3D kinematics of end effector trajectories to understand the nominal cyclic leg trajectories and the variability around those kinematic cycles (Fig. 1c). We hypothesize that the combination of body dynamics and attachment strategies generates an inverted running gait that is highly robust, allowing fast feed-forward running, despite considerable variability.

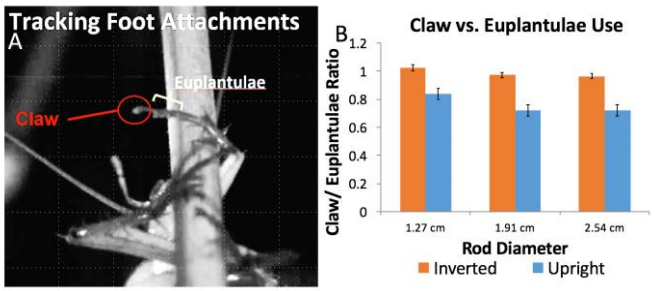


Fig. 11. Branch/pole running in cockroaches. A) Photo showing the claw and euplantulae or friction pads during inverted pole running. B) Claw/euplantulae ratio equals the stance time of claw ÷ stance time of euplantulae. Cockroaches used their claws more during inverted running.

b) Probabilistic graphical models of running variability and stability

We empirically discovered the relationship between the frequency of transition to inverted running and rod diameter. The incredibly high kinematic variability across trials suggests modeling this behavior including stochastic dynamics. Modeling cockroach branch running using probabilistic graphical models (e.g. Markov chain) allows us to not only capture state dependent variability, but also the rod diameter dependent stability. As the rod diameter increased the expected amount of time the cockroach can run on the rod before transitioning into an inverted posture increases. In future work, we will incorporate motor control into this modeling framework to predict how different controllers affect stability in the face of high kinematic variability.

C. Characterizing Multi-component, 3D Terrain Using Photogrammetry

1) Accurate models of 3D terrains and insect morphology

In our previous experiments, we have pieced together components of 3D terrain to represent nature sufficiently. However, these were single component studies and we propose to study more complex, multi-component terrain in the future. With the advancement of 3D image capture techniques, we propose to explore photogrammetry to create 3D models of these terrains. This technique creates 3D meshes of an object from few images of the object. Thus, photogrammetry is a simple, accessible, and mobile technique.

We began by obtaining 3D meshes of a multi-component terrain reconstructed in the laboratory. As a first step, we imaged an individual plant and potting soil surrounding it (Fig. 12, Middle) as a sample of multi-component terrain. We followed with imaging an outdoor patch of grass (Fig. 12, Top).

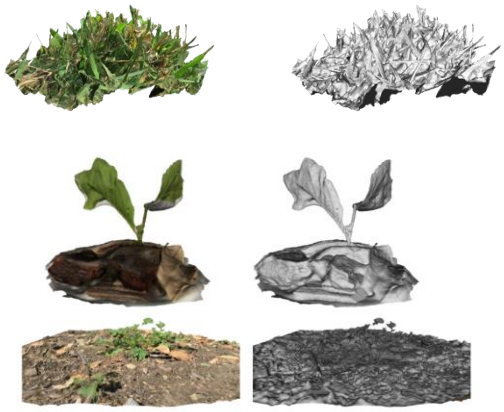


Fig. 12. Preliminary efforts at photogrammetry to characterize 3D, multicomponent terrains. Outdoor grass patch (Top). Indoor plant (Middle), and outdoor terrain (Bottom). Color images are on the left with their corresponding mesh on the right.

We have also reconstructed a 3D model of one cockroach species - *Blaberus discoidalis* which we used for our multi-component terrain studies previously (Fig. 13). In addition, we have reconstructed a goliath bird-eater tarantula, giant desert scorpion, a flower beetle and grasshopper.

These meshes can be exported in nearly all common mesh formats compatible with other CAD software. We propose to use these 3D models to simulate the movements of insects through complex terrains and navigation through confined spaces. By running these simulations with insects of diverse body and limb design, we aim to understand what morphological features permit or constrain their locomotion through these terrains. Based on this, the geometry of the robots we build will be modified, enhanced, and optimized.

Even though photogrammetry works irrespective of the size of a sample in principle, it is a challenge to obtain high resolution images of small-sized objects without any blur in the images. Using macro lenses meant that we had low depth of field and some part of the object was not in focus. This issue

can be resolved by taking a greater number of photographs to focus on smaller portions, and then stitching these photographs to obtain a single image with all parts in focus (focus stacking). This would mean that the time needed to take photographs for a single sample would increase significantly. Another challenge is the inability to get diffused lighting and avoid any reflection from the surface, such as produced by the insect exoskeleton while obtaining images. We have begun addressing both these shortcomings by building a photo studio for appropriately sized objects with multiple diffused light sources and multiple cameras to take many photographs.

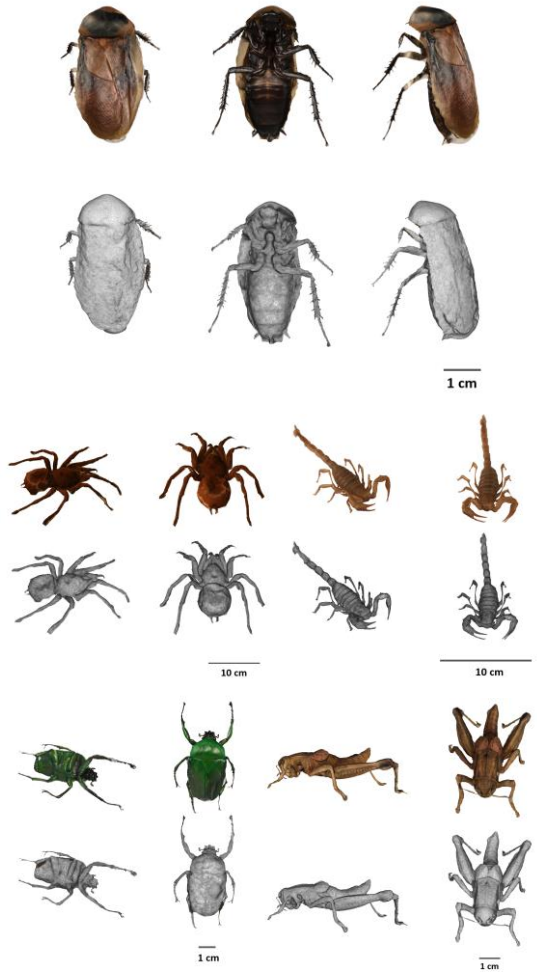


Figure 13. Preliminary efforts at photogrammetry to characterize diverse animal morphologies. Color images are on the top with their corresponding meshed below. Shown in order are a discoid cockroach, goliath bird-eater tarantula, giant desert scorpion, a flower beetle and grasshopper

To extend using photogrammetry to model real-life multi-component 3D environment, we have attempted to capture outdoor terrains. We chose a grass patch and obtained a 3D mesh (Fig. 12, Top). The software (Autodesk Remake) allows the user to set a scale by defining a known distance in the 3D mesh. After calibration, we measured the geometric properties of features of interest on the 3D mesh and the distribution of these features in the terrain.

As a first step towards characterizing geometry, we measured the heights of the grass blades and distances between them. The distance between grass blades represents path widths available for the cockroaches to navigate through. The grass blade height in our model was measured to be 6.65 $\pm$ 2.28 cm (mean $\pm$ s.d.). The path widths ranged from 2 cm to 7 cm (4 $\pm$ 1.33 cm). The grass density was measured to be approximately 238 blades/m^2.

We also obtained a 3D mesh of a *Blaberus discoidalis* using the Next Engine 3D laser scanner to contrast with photogrammetry. On comparing the scans obtained from both, we observe that the laser scan has much lesser feature resolution. The body also has very small features, like the spines on their feet. The cockroach body has a lot of occlusions, especially on the ventral side where the limbs attach to the body. Combined with the limited range of roll of the turntable the resultant scans might have a lower resolution in these occluded areas. Photogrammetry is at an advantage here. Increasing the number of photographs and capturing high occlusion regions from different angles enables the user to achieve higher resolution in these regions. The use of macro lenses also allows for the capture of very fine features in high resolution. A photogrammetry scan if done using focus stacking techniques takes about 5 hours to capture photographs of a single sample. The laser scanner on the other hand only requires the user to set the sample on the turntable and input the capture parameters in the software and scanning is automatic. This can however be overcome by building an automated rig in the future for photogrammetry. In conclusion, we have chosen to use photogrammetry for our further research experiments because of the higher fidelity of morphological reconstruction.

2) Navigating through natural terrains

In a real environment, the outdoor terrain can include grass patches, rocky areas, areas with wooden sticks, dried leaves, and more generally a combination of these. Since the robots must navigate through these complex terrains, it is critical to be able to characterize the terrains to determine the most effective path as well as its control strategy through the path. We aim to do this by building a cloud-based database of terrains, which the robot on the ground could access to help determine its locomotion strategy. We have begun the data collection for obtaining 3D models for each of these terrain types to add to the database. Once the robot has access to the 3D model of the terrain, it can assess the trade-offs to choosing a particular path enroute to a goal. We are currently developing a set of metrics to characterize and classify terrains.

Given how different terrains can be, we plan to divide the characterization into steps. Our first step would be to use the 3D scanned models to assist in identifying the type of the terrain - grass, rocks, dried leaves, concrete debris, wooden twigs, etc. We aim to do this by having a large database with sample 3D models with pre-determined terrain types. At some

point in the future, we imagine that a robot could compare its terrain to those stored in a cloud's database to identify its type.

The second step is to gather more information concerning the properties of the terrain both geometrical and material properties. For example, in case of a grass-like terrain, we would add material properties like the stiffness and roughness of the blades. Stiffness would determine the degree of bending of the grass blades as the robot traverses them. The roughness of the blades could affect the contact forces the robot experiences. In addition, we would add a determination of the properties of the substrata. To develop such a resource database, we need to measure and catalog properties for each identified type of terrain that the robot is likely encounter for a given task. We have collected some material property data for different grass blades using a materials testing machine (Instron). These above data are assisting in the design of our laboratory testbeds, but will begin to be used to decide effective navigation paths in the future.

Admittedly, the dimensionality of all measurable terrain features is very large. However, by making laboratory performance tests, we hope to learn if we can effectively collapse dimensions to produce reduced order models that might give us sufficient information to select preferable paths based on potential performance trade-offs. We hypothesize that there exist signature properties that might be most influential across different terrains, and hence reduce the data and time needed by controllers for decision making.

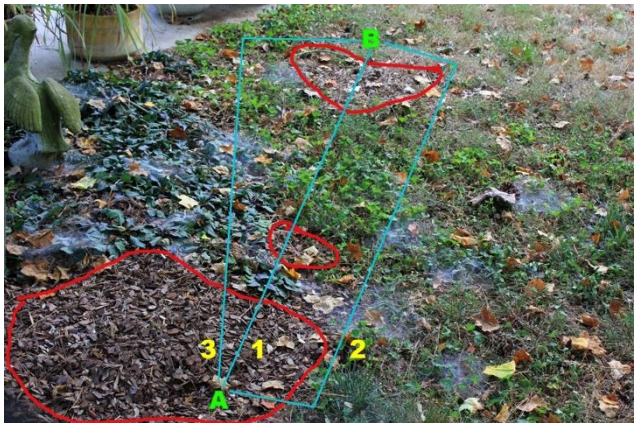


Figure 14. Path selection in natural terrain. Preliminary efforts at photogrammetry to characterize the trade-offs of possible paths in the environment. Potential paths are numbered and challenging areas are circled in red.

The final goal of using our terrain database would be to assist a robot in choosing its path. While choosing a path, the robot will undoubtedly have multiple performance goals to optimize, including time, path where it would need to maneuver least, lowest energy, path of least failure probability (e.g. avoid paths where are regions where the robot might get damaged or stuck) etc.

Consider a ground terrain example (Fig. 14). Assume a robot is required to navigate from point A to point B. Based on the data from a cloud terrain database, we could identify zones (marked in red) of wooden debris, small green shrubs, and dried leaves. The robot could select among multiple path options (drawn in blue). Path 1 would be the shortest, but the robot must navigate through three different challenging terrains. This path would have the highest probability of failure. Path 2 would require it to traverse the least wooden debris and then travel through grass, avoiding the dried leaves patches. The length of this path would be the longest and the time likely the greatest. Path 3 would avoid more dried leaves and only travel through wooden debris and green leaves. The robot would need to maneuver the most while it ventures through the leaves. The time needed for traversal will be determined by the 3D path length and properties of the terrain. The decision criteria for paths could be tabulated as in Table I. Based on such tradeoffs, and the requirements of the task, the robot could be continually updated to choose the optimal path.

TABLE I
TERRADYNAMIC DECISION PATHS

	Path 1	***Path 2***	***Path 3***
Time	Low	High	Medium
Need for maneuvering	Medium	Medium	High
Probability of Failure	High	Low	Low

V. Conclusions & Future Work

Clearly multi-component 3D terrain is at least as complex as aero- and hydro-dynamics in the real environment. Yet, animals manage to negotiate terrain with the ease and effectiveness unavailable to current robots. Here, we detail continued progress at how animals to squeeze through cracks and maintain locomotion in confined environments, likely use sensory feedback to navigate grass-like beams, and use multi-component feet to traverse branches/rods at high speed.

We support the conclusions of our Berkeley colleagues [12] when they state that "The Cloud infrastructure and its extensive set of Internet-accessible resources has potential to provide significant benefits to robots and automation systems. We consider robots and automation systems that rely on data or code from a network to support their operation, i.e., where not all sensing, computation, and memory is integrated into a standalone system." Here we began to use scanning technology to not only scan the bodies of animals that serve as our inspiration, but also complex 3D multi-component terrain.

At present, photogrammetry has considerable limitations, but we are convinced that as progress is made to improve the images that these will provide a foundation for a cloud database of terrain. The terrain database will be helpful with respect to the general guidance of ground robots with respect to identifying areas of effective traversal and those more likely to serve as challenging barriers to travel or ingress/egress. To complement the 3D geometrical meshes of terrain, we could add the material and surface properties of the objects and substrata of the natural environment. The combination of

geometry and material properties could allow remarkably detailed terrain path suggestions for mobile ground robots.

VI. References

[1] C, Li, Z. Zhang, and D. I. Goldman. "A terradynamics of legged locomotion on granular media." *Science*. 339, 1408, 2013.

[2] T. Libby, T.Y. Moore, E. Chang-Siu, D. Li, D. J. Cohen, A. Jusufi, and R. J. Full. "Tail-assisted pitch control in lizards, robots and dinosaurs." *Nature*, *481*(7380), 181-184, 2012.

[3] S. Sponberg and R. J. Full. "Neuromechanical response of musculo-skeletal structures in cockroaches during rapid running on rough terrain". *Journal of Experimental Biology*. 211, 433-446, 2008.

[4] J. C. Spagna, D. I. Goldman, P-C. Lin, D. E. Koditschek and R. J. Full. "Distributed mechanical feedback in arthropods and robots simplifies control of rapid running on challenging terrain". Bioinsp. Biomim. 2, 9–18, 2007.

[5] A. J. Spence, S. Revzen, J. Seipel, C. Mullens, K. Yeats, and R. J. Full. "Insects running on elastic surfaces". *Journal of Experimental Biology*. 213, 1907-1920, 2010.

[6] J. M. Mongeau, B. McRae, A. Jusufi, P. Birkmeyer, A. M. Hoover, R. Fearing, R., and R.J. Full. "Rapid inversion: Running animals and robots swing like a pendulum under ledges". PloS one, 7(6), e38003, 2012.

[7] F. Qian, and D. Goldman. "Anticipatory control using substrate manipulation enables trajectory control of legged locomotion on heterogeneous granular media". In *SPIE Defense+ Security* (pp. 94671U-94671U). International Society for Optics and Photonics, 2015, May.

[8] B. McInroe, H. C. Astley, C. Gong, S. M. Kawano, P. E. Schiebel, J. M. Rieser, H. Choset, R. W. Blob and D. I. Goldman. "Tail use improves performance on soft substrates in models of early vertebrate land locomotors". *Science*, 353(6295), pp.154-158, 2016.

[9] K. Jayaram, and R. J. Full. "Cockroaches traverse crevices, crawl rapidly in confined spaces, and inspire a soft, legged robot." *Proceedings of the National Academy of Sciences*, *113*(8), E950-E957, 2016.

[10] C. Li, A. O. Pullin, D. W. Haldane, H. K. Lam, R. S. Fearing, and R. J. Full, "Terradynamically streamlined shapes in animals and robots enhance traversability through densely cluttered terrain," *Bioinspir. Biomim.*, vol. 10, no. 4, p. 46003, 2015.

[11] N. H. Hunt, H. S. Moon, A. Hammond, N.P. Burnett, M. Pritchard-Berman and R.J. Full. "Stability of cockroaches running rapidly on rigid rods." In *Integrative and Comparative Biology* (vol. 54, pp. e97-e97). Oxford Univ. Press Inc., 2014, January.

[12] B. Kehoe, S. Patil, P. Abbeel, and K. Goldberg. "A survey of research on cloud robotics and automation." IEEE Transactions on Automation Science and Engineering, 12(2), 398-409, 2015.

Bio-inspired Principles of Transitions

MAST MCE-17-2.3

Robert J. Full, Chen Li, and Thomas Libby
Department of Integrative Biology
University of California at Berkeley
Berkeley, CA, USA

Chad C. Kessens
United States Army Research Laboratory
Aberdeen Proving Ground, MD 21005, USA

***Abstract*—Small MAST-scale terrestrial robots have achieved remarkable speed and versatility, but still lack the agility and robustness of animals. Here we describe efforts to understand two locomotor transitions – self-righting and horizontal plane maneuvering. Inspired by the multi-functionality of the discoid cockroach's wings (closed for obstacle traversal, opened for self-righting), we designed a novel winged self-righting mechanism based on our recently-developed rounded shell for obstacle traversal to further enable dynamic self-righting. By systematic experiments and modeling, we discovered mechanical principles of winged self-righting to guide the design and control of self-righting wings. To expand our understanding of the possible mechanisms of self-righting, we began study of righting in small lizards. Lizards showed larger diversity in approach than individual cockroach species, and appear to use inertial forces from their tail as in air-righting. Our study of planar maneuverability revealed surprising features of unsteady locomotion, as lizards used highly variable (but strikingly robust) limb function to execute consistent, rapid transition performance. Inspired by the form and function of the lizards limbs during turning, we created the most rapid-turning legged robot yet reported.**

***Index Terms*—bio-inspiration, transitions, self-righting, turning, legged robots**

I. Introduction

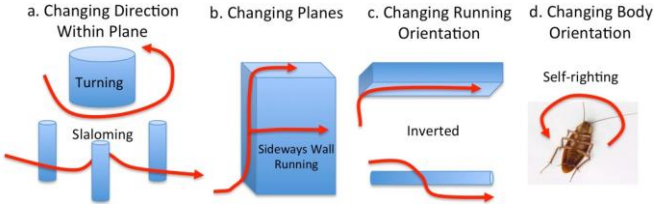


Fig. 1. Transitions. Animals in nature and MAST scale robots (red line paths) must negotiate a) turns and obstacles, b) wall transitions, c) inverted surfaces - ceilings, rods, and d) recover from being flipped upside-down.

A. Background and Technical Challenges

MAST Ground Platforms are becoming increasingly effective at running and climbing, even through granular media (MCE-2.1) and multicomponent 3D terrain (MCE-2.2). Yet, these platforms will never meet task objectives in the field unless they can execute transitions that include agile changes in direction within a plane, from the ground to a wall or ceiling and self-righting from an inverted position (Fig. 1). Fortunately, animals negotiate these transitions often with ease and at high tempos. Here, we propose to discover the **fundamental biological principles of transitions** to inspire the design of the next generation of MAST Ground Platforms, so as to further develop their unique capabilities at this size scale.

B. Research Objectives

At least two major challenges must be met to quantify and model transitions. We need to explore and categorize behavioral strategies relevant to translation to robots. From these categories, we will **quantify the complete dynamics of transitioning from one stable behavior to another** so that the most effective strategies are selected with respect to agility, time, energy, and simple control. To understand the mechanisms used, we need to conduct experiments on the animals to determine which morphological features permit the novel capabilities of changing planes and orientations. This requires manipulating feet, legs, bodies, sensors and surfaces. The second challenge is the ability to **create reduced order models of dynamic transitions** representing the variety of behaviors that small animals display, so we can extract general biological principles translatable to MAST robots. We hypothesize that the exceptional, and seemingly unstoppable, locomotor performance of animals will be attained by: (1) using existing appendages in novel ways to execute the transition (e.g. wings to self-right), (2) using alternate mechanisms that rely on interaction of the body and appendages with environment to adapt current strategies or generate new behaviors, and (3) using dynamic changes in body/tail shape to support or augment impulses from legs. We performed experiments that simulate relevant transitions to elucidate novel, high-tempo locomotor behavior in animals.

II. Related Work

Synergy with Existing Works & External Collaborations

1. Synergy with Task MCE-2.2. In Task 2.2, we enabled the VelociRoACH robot to traverse cluttered obstacles by adding a "terradynamically streamlined" rounded ellipsoidal shell (Fig. 2c). However, the robot still lacked the animal's multi-functional abilities [1]: when it over-rolled during

obstacle traversal, the robot became permanently flipped over because the rounded shell is stable when upside down (Fig. 2c) [2]. Here, we further enabled the robot to self-right by modifying the shell into actuated wings and studied the principles of winged self-righting as a next step towards multi-functional locomotion.

2. Collaboration with Kessens at ARL. To leverage existing expertise at Army Research Lab on self-righting modeling, we collaborated with Chad Kessens and utilized a generic planar self-righting analysis framework [3] which he developed.

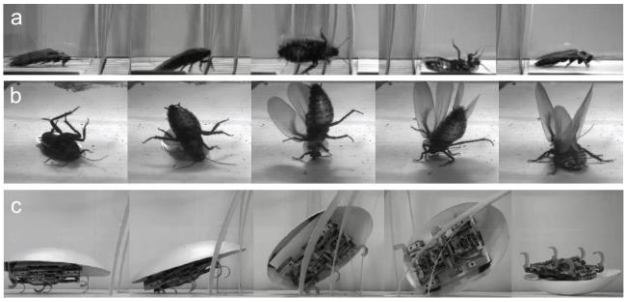


Fig. 2. Shell use for transitions. (a) A discoid cockroach rapidly traverses cluttered obstacles such as grass-like beams [2], flips over, and quickly rights itself. (b) The cockroach self-rights by opening and pushing its wings against the ground [4]. (c) A small legged robot uses a cockroach-inspired rounded shell to traverse cluttered obstacles; however, when it over-rolls and flips over, it becomes stable and cannot self-right [2].

III. Technical Approach

A. Biological Inspiration for Robot Self-righting

Small terrestrial animals and robots are both susceptible to flipping-over during rapid locomotion in complex terrains. However, small robots are much less capable of self-righting from an upside-down orientation compared to small animals like insects.

To self-right after flipping-over, many small animals use exaptations [5] of appendages that primarily serve other purposes. For example, the discoid cockroach that lives in cluttered forest floor foliage has wings that are normally folded against the body to form a protective "shell". In Task 2.2, we have discovered that this shell helps the animal traverse cluttered obstacles such as grass-like beams, because its rounded shape is "terradynamically streamlined" and reduces terrain resistance by facilitating body rolling [2] (Fig. 2a). However, when flipped over (Fig. 2a), the discoid cockroach can also open its wings and use them to rapidly push against the ground to right itself (Fig. 2b) [4].

B. Geometric modeling for Insect Self-righting

A central hypothesis that has emerged from our animal experiments is that animals more often self-right by using body rotations that overcome lower potential energy barriers [4]. To test this hypothesis, we developed a simple 3-D geometric model to estimate the animal's maximal center of mass (CoM) height (Fig. 3). We approximated the body shape of each of three species of cockroaches as a rigid ellipsoid with length 2*a*, width 2*b*, and thickness 2*c* from its morphological measurements (Fig. 3a-c), approximated the CoM position with geometric center, and assumed that the lowest point of the rigid body always touched the ground.

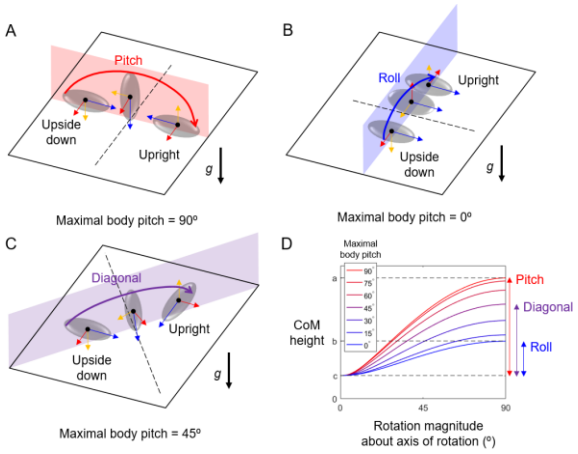


Fig. 3. Schematics of a rigid ellipsoid approximating the animal body, which rotates about the pitch (a), roll (b), and diagonal (c) axes that lie within the horizontal ground plane. (d) Center of mass height of the rigid ellipsoid calculated from the model during the rising phase as a function of rotation angle in the vertical plane of rotation for seven representative maximal body pitch angles (values shown in legend).

We carefully observed high-speed videos and approximated the animal's body rotation with the ellipsoid rotating about a fixed axis that lied within the horizontal ground plane. We then measured the maximal body pitch, θ_{pitch_max}, the angle between the long axis of the ellipsoid and the horizontal ground plane when the CoM reached its highest position at the end of the rising phase. Our simple model showed that maximal CoM height and thus at the end of the rising phase increases monotonically with maximal body pitch, and so does the potential energy barrier that the body has to overcome (Fig. 3d). Therefore, we carefully examined high-speed videos of animal self-righting and measured the maximal body pitch for three species of cockroaches during self-righting, and used use this model to estimate the potential energy barrier heights for animal self-righting trials.

C. Geometric Modeling of Robot Self-righting

To understand the energetic requirements for winged self-righting and to determine whether the robot and the discoid cockroach could self-right quasi-statically or must right dynamically, we applied the geometric modeling framework [3] that Kessens developed. We approximated the body shape in the sagittal plane as a rectangle and the wing shape in the sagittal plane as a rigid truncated ellipse (Fig. 4a). To study how wing shape affects self-righting, we fixed the wing length and the truncation distance while allowing the wing height ϵL to vary, where ϵ is wing height normalized to wing length (Fig. 4b). A small ϵ yielded thinner wings, while a larger ϵ yielded taller wings. To study how the mass distribution between the body and the wings affected self-righting, we also varied relative wing mass $\mu = m_{wings}/(m_{body} + m_{wings})$, the ratio of wing mass to total mass.

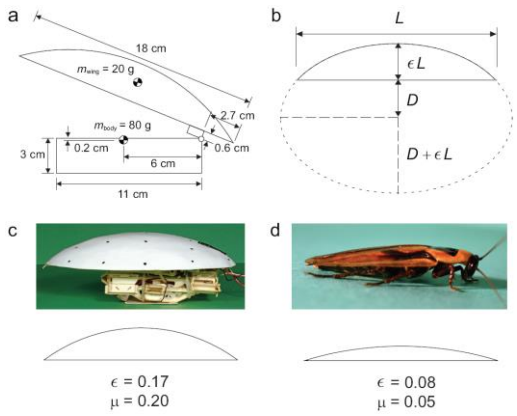


Fig. 4. Geometric model. (a) Schematic of the body as a rectangle and wings as a truncated ellipse in the sagittal plane. (b) Parameterization of the truncated ellipse. (c, d) Side views of the robot and the discoid cockroach.

D. Self-righting in Lizards

To determine and implement the most effective self-righting strategy to MAST Platforms, we must have a better understanding of principles of self-righting in biological systems. To achieve this, we began to broaden our comparative approach to include lizards. House geckos, *Hemidactylus platyurus*, have a rapid air-righting response [6] which harnesses inertial forces generated by tail swings. This approach, if used on the ground, would represent a physically distinct mechanism from those discovered in invertebrates. We used an eight-camera tracking system in to observe 3D righting kinematics in high speed. To prevent body markers being lost when the animal is inverted, we used a glass plate substrate to enable views from below.

E. Horizontal Plane Agility in Lizards

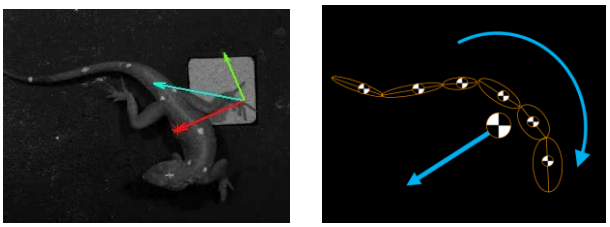


Fig. 5. Lizard horizontal plane maneuver. Left: Agamid lizard during horizontal plane maneuver, with single-leg ground reaction forces measured by miniature force plate. Right: rigid body model informed by animal morphometrics estimates body CoM and linear and angular momentum about CoM.

We used prey capture events to elicit near-maximal effort planar maneuvers in lizards, resulting in a well-defined task: start from rest, turn 180°, and accelerate towards the prey as quickly as possible. We video recorded manevuers at 500 fps and extracted kinematics by digitizing points along the back of the animal. We estimated angular momentum about the animal's total center of mass CoM at each frame in the video using a six-link rigid body model (two segments for the body, and four for the tail; Fig. 5 Right). This required an extensive morphometric model with measurements of length, mass, segment CoM and moment of inertia of each link. Morphometrics for living animals were estimated from averages from cadavers.

F. Horizontal Plane Agility in Robots

We complemented our animal study of the principles of horizontal plane agility by building two generations of a minimally-actuated physical model (robot; Fig. 6). In our first effort, we hewed close to the design of VelociRoACH, hypothesizing that we could implement fast turning with just two DoF – a lateral-only DoF in the front, and a single DoF driving a pair of hind limbs 180° out of phase in the rear (as in VelociRoACH's steady running condition). Based on preliminary experiments, we added an additional DoF in the rear, supporting independent rotation of the hind limbs.

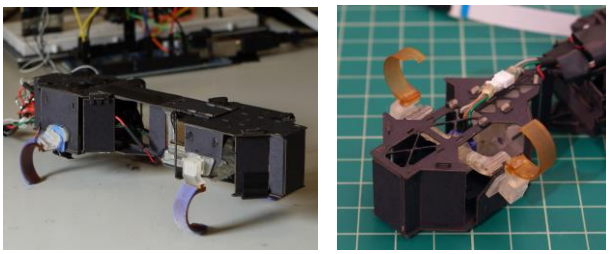


Fig. 6. Turnbot prototype. Left: 3 DOF turning robot with independent sagittal plane rear limbs. Right: Closeup of front limbs with single lateral DOF.

IV. Results

A. Metrics for Performance Evaluation

Research accomplishment summary 2016. Developed robots and elucidated principles of winged self-righting and legged turning. Revealed function of limbs during agile planar maneuverability task in sprawled runners. Began study of tailed terrestrial self-righting.

Plans for 2017. Integrate lessons from turning and righting robots into platforms with capability for running on rough terrain. Continue study of tailed terrestrial self-righting.

B. Insect Self-righting: Potential Energy Barriers

Using the measured maximal body pitch (Fig. 7a-c), we obtained histogram of the potential energy barrier height (Fig. 7d-f) from our simple geometric model. We found that all three species had higher probabilities of overcoming lower potential energy barriers during successful self-righting maneuvers. The potential energy barriers that the Madagascar hissing cockroach and the American cockroach overcame during successful self-righting maneuvers was almost always near the lower limit ($b - c$), which was for pure body rolling (Fig. 7d, e). The Madagascar hissing cockroach never overcame high barriers of heights near the upper limit ($a - c$), which was for pure body pitching. The American cockroach only did this once out of over 50 trials, during the only trial in which it used wings to right. By contrast, the discoid cockroach overcame potential energy barriers of a broader range of heights, from around that for pure body rolling ($b - c$) to near that for pure body pitching ($a - c$) (Fig. 7f). Despite

the broad distribution, there was a clear anti-correlation of the probabilities with the heights of potential energy barriers.

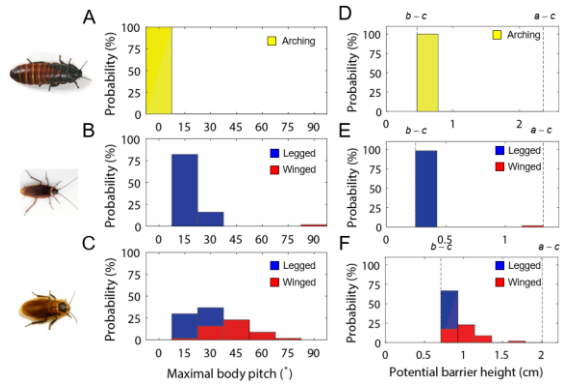


Fig. 7. (a-c) Histogram of the maximal body pitch observed during the rising phase of self-righting, and (d-f) histogram of potential energy barrier height estimated from the simple rotating ellipsoidal rigid body model for the three species of cockroaches.

Overall, our experimental observations and model calculations confirmed our hypothesis that the animals indeed more often self-right by using body rotations that overcome lower potential energy barriers [4].

C. Self-righting Robot Design & Development

Inspired by the winged discoid cockroach, we designed a new robot that opens its wings to self-right by pushing against the ground (Fig. 8), and used this robot to systematically study the principles of winged dynamic self-righting [7]. With future integration of self-righting and obstacle traversal [2] capabilities in mind, we formed the two actuated wings by sagittally slicing the same rounded ellipsoidal shell previously developed for obstacle traversal, adding actuators and transmissions for each half (Fig. 8a, b).

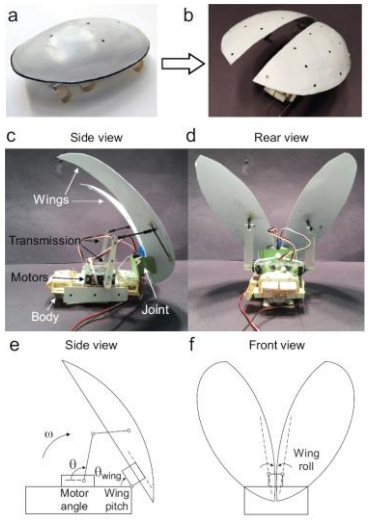


Fig. 8. Self-righting shell design. (a, b) The "terradynamically streamlined" rounded shell previously developed for obstacle traversal was transformed into two actuated wings for self-righting. (c, d) Side and rear views of the actuated wings mounted on a VelociRoACH robot body. (e, f) Side and front view design diagrams showing the pitching and rolling degree of freedom of the wings.

When folded against the body, the two wings formed a rounded shell, similar to the discoid cockroach (Fig. 8b). We designed a novel two degree-of-freedom (DoF) joint and four-bar transmissions to allow the wings to open in a similar fashion to the discoid cockroach (Fig. 8c, d). The wings were attached to the anterior end of the robot body *via* the 2-DoF joint, with one degree of freedom for pitching motion (Fig. 8e) and the other for rolling motion (Fig. 8f) of the wings relative to the body. The two wings were actuated by two small, lightweight, high-torque servo motors, and could be either opened in exactly the same way (symmetric wing opening) or differently (asymmetric wing opening).

D. Robot Self-righting using Symmetric Wing Opening

To begin to discover the principles of terrestrial self-righting using wings, we performed robot experiments on a level, flat, rigid ground by controlling the two wings to open in synchrony (symmetric wing opening) with the same magnitude and speed. Similar to the discoid cockroach, as the robot opened and pushed its wings against the ground, its body pitched up, rotated about the ground contact, and eventually vaulted over the front edge of the wings and fell to the ground in an upright orientation, resulting in successfully righting (Fig. 9a, b; Fig. 9c, red region). Failed righting maneuvers resulted in the body settling into a nearly vertical stationary orientation (Fig. 9c, blue region).

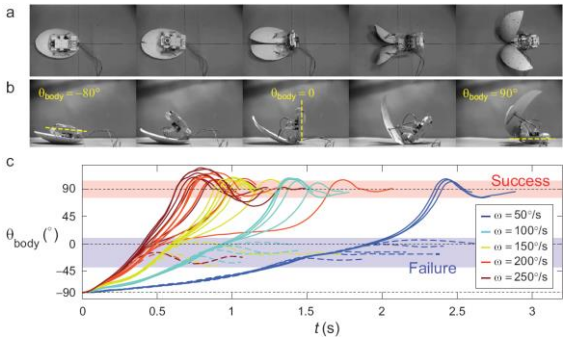


Fig. 9. Symmetric wing opening. (a, b) Top and side views of the robot during a successful self-righting maneuver at a few representative instances, with definition of body pitch angle, θ_{body}. (c) θ_{body} as a function of time, t, for a wide range of motor speeds ω and angle θ tested. Each curve of the same color uses the same ω (see legend), but different θ. Solid and dashed curves represent successful and failed maneuvers, respectively.

E. Effect of Wing Opening Magnitude & Speed on Righting Performance

To understand how wing opening affects the robot's self-righting performance, we systematically tested how wing opening magnitude and speed affected the robot's performance (Fig. 10), including righting probability, P_{right} (success = 1, failure = 0), and righting time, t_{right}.

The robot's self-righting probability increased with both wing opening magnitude and wing opening speed (Fig. 10a). For large wing opening magnitudes ($\theta > 110°$), the robot always righted ($P_{right} = 1$), even when wing opening was the slowest tested ($\omega = 50°/s$). For small wing opening magnitudes ($\theta < 88°$), the robot always failed to right ($P_{right} = 0$), even

when wing opening was the fastest tested (ω = 250°/s). For intermediate wing opening (88° ≤ θ ≤ 110°), the robot righted when wing opening was faster (larger ω), but failed when wing opening was slower (smaller ω). In addition, the robot self-righted more quickly as wing opening became larger or faster (Fig. 10b): t_{right} decreased with ω for any given θ and decreased with θ for any given ω.

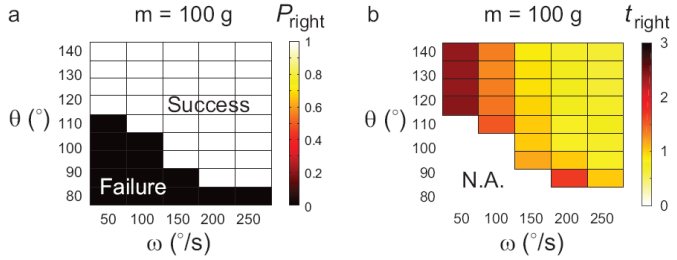


Fig. 10. Righting performance. (a) Righting probability, P_{right} (success = 1, failure = 0), and (b) righting time, t_{right}, as a function of motor angle θ and motor speed ω. Brighter colors: higher performance; darker colors: lower performance.

F. Effect of Wing Shape & Mass on Quasistatic Self-rightability

To understand how wing shape and mass distribution affect the kinematic requirement for dynamic self-righting, we varied relative wing height and relative wing mass and studied how they affected the minimum wing pitch angle required, θ_{req}^{wing}, for quasi-static self-righting (Fig. 11a). We found that, as long as the wings weighed less than the body (μ < 0.5), θ_{req}^{wing} always decreased as the wings became taller (increasing ϵ). For any given relative wing height (given ϵ), θ_{wing}^{req} was always larger when the wings were heavier (larger μ). When the wings became as heavy as or heavier than the body (μ ≥ 0.5), the robot was unable to right itself quasi-statically, because the lighter body lifted off the ground instead of the heavier wing. Together, these observations demonstrated that a robot (or animal) with taller and lighter wings can more easily self-right quasi-statically, which may be necessary if it is incapable of opening wings quickly (for example, if the wing joints were highly geared).

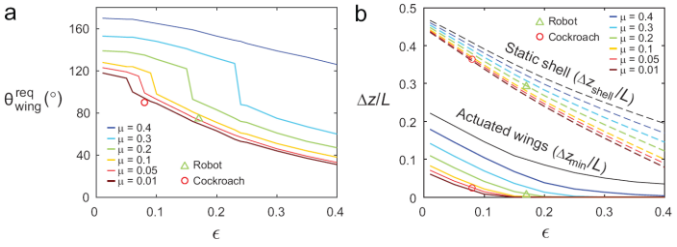


Fig. 11. Quasistatic self-rightability (a) Minimal wing pitch angle required, θ_{wing}^{req}, for quasi-static self-righting as a function of normalized wing height ϵ for a wide range of relative wing mass μ. (b) Minimum potential energy barrier height Δz (normalized by wing length *L*) that the robot must overcome to self-right from the stable overturned orientation with wings closed as a function of ϵ for a wide range of μ. Solid: if the robot can open wings to θ_{wing}^{max} = 75°. Dashed: if the wings cannot be opened, i.e., if they are a static shell. In (a) and (b), green triangle and red circle represent model predictions for the physical robot and the discoid cockroach, respectively.

For both the robot (ϵ = 0.17, μ = 0.20) and the discoid cockroach (ϵ = 0.08, μ = 0.05), the model predicted that it must open its wings to $\theta_{wing}^{req} \approx 90°$ (Fig. 12), beyond the maximal wing pitch angle that the robot and the discoid cockroach are capable of (Fig. 11a, green triangle and red circle). This means that the robot and the discoid cockroach could not self-right quasi-statically and must use kinetic energy to right dynamically.

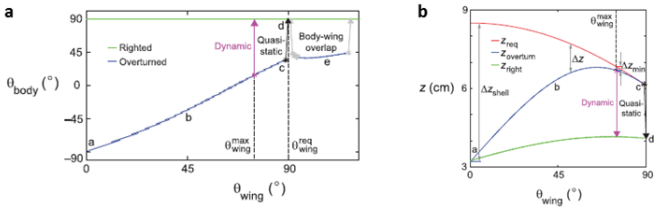


Fig. 12. a) State space map showing how overturned states (blue) can reach the goal states (green). b) Center of mass height for static wing (shell), overturned states as a function of wing angle, and minimum required height to achieve self-righting (req). The difference between z_overturn and z_req is the potential energy barrier that must be overcome for a given wing angle.

G. Effect of Wing Shape & Mass on Potential Energy Barrier

To understand how wing shape and mass distribution affect the energetic requirement for dynamic self-righting, we determined the minimal potential energy barrier height, Δz_{min}, that must be overcome by using kinetic energy for the wide range of relative wing height and relative wing mass tested (Fig. 11b, solid curves). This value was dependent on the range of wing angles (Fig. 12b). Within the wing opening that the animal and robot were capable of (0 ≤ θ_{wing} ≤ 75°), we found that, for any given μ, Δz_{min} decreased monotonically with ϵ; for any given ϵ, Δz_{min} was always higher for larger μ. Therefore, tall, lightweight wings are energetically advantageous for winged dynamic self-righting.

To study how much energy winged self-righting helped the robot and animal save compared to if they do not use wings (such as turtles [8]), we further calculated the potential energy barrier height, Δz_{shell}, with a static shell (θ_{wing}^{max} = 0) (Fig. 11b, dashed curves). We found that Δz_{shell} was much higher for a static shell than for using wings to self-right: for example, by a factor of 2 for μ = 0.5 and an order of magnitude for μ = 0.05. Therefore, winged dynamic self-righting is more economical than passive self-righting using rigid shells (e.g., turtles with highly-domed shells [8]).

H. Robot Self-righting using Asymmetric Wing Opening

In our animal experiments [4], we hypothesized that it may be advantageous to use asymmetric wing opening to self-right. This was based on the observation that the animal rarely opened the two wings symmetrically and righted by pitching over its head in the sagittal plane, but instead usually opened the two wings asymmetrically and righted using body rotations out of the sagittal plane.

To test this hypothesis, we tested our robot by opening the left and right wings by to different motor angles, θ_L and θ_R. We observed that, with asymmetric wing opening, the robot body rotation was not purely pitching within the sagittal plane,

but also had a rolling component. This led to more diverse righting outcomes than using symmetric wing opening 13a).

To examine whether asymmetric wing opening enhanced the probability of the robot self-righting, we measured average righting probability for each combination of θ_L and θ_R tested. We found that for large θ_L and θ_R (> 100º), P_{right} was usually high (> 80%) and relatively less sensitive to wing opening asymmetry (13b, top right quarter). Surprisingly, for small θ_L and θ_R (≤ 90º), asymmetric wing opening consistently resulted in higher P_{right} (≈ 50%) than symmetric wing opening did (≈ 0) (Figure 12b, bottom left quarter). This suggests that when wing opening magnitudes are limited (e.g., when an animal fatigues or when a robot has low power), asymmetric wing opening increases the chance of righting.

Overall, our systematic experiments using the robot as a physical model revealed the mechanical principles of winged dynamic self-righting and advanced our understanding of how to design and control wings for small legged robots to self-right rapidly and economically.

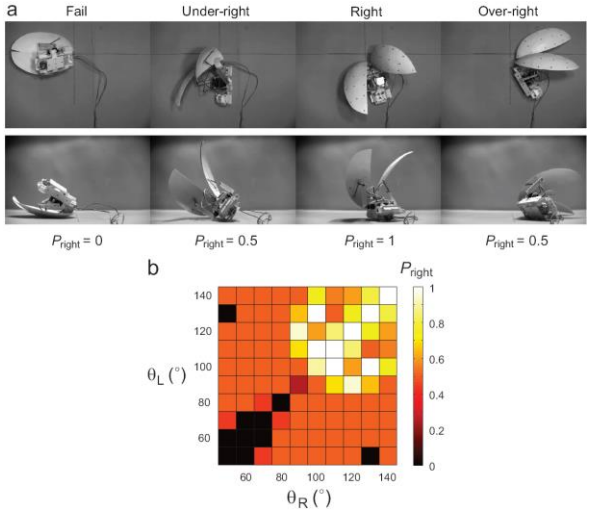


Fig. 13. Asymmetric wing opening. (a) Top and side views of the robot's final body orientation for the four outcomes observed: failure, under-righting, successful righting, and over-righting. (b) Average righting probability, P_{right}, as a function of left and right wing opening magnitudes, θ_L and θ_R.

To estimate the potential energy barrier that the robot overcame during self-righting using asymmetric wing opening, we developed a simple geometric model of the robot. Based on experimental observations, we treated the robot as an ellipsoid with uniform mass distribution rotating about a fixed axis within the ground plane (Figure 14a), (i.e., center of mass was at geometric center). Using this simple geometric model, we varied the angle between the rotation axis and the initial fore-aft axis of the body (β /2), and numerically calculated the potential energy barrier height (the increase of center of mass height during the rising phase), Δz, as a function of β, assuming that the lowest point of the body touched the ground (Figure 14b). Using the model-predicted dependence of center of mass height on change of heading and the average change of heading measured from the asymmetric righting experiments, we then estimated the center of mass height z_{max} for all the θ_{wing}^L and θ_{wing}^R tested (Figure 14c).

We found that Δz decreased with the lower of the two the wing opening magnitudes and was lowest when the robot righted using asymmetric, small wing opening. This suggests that, besides increasing righting probability, asymmetric wing opening also saves energy by reducing the potential energy barrier height when wing opening magnitudes are limited. Our asymmetric righting experiments and modeling together suggest that the observed frequent occurrences of asymmetric wing opening may be an adaptation of the discoid cockroach to right more quickly and economically, and can be a useful strategy for robots to increase self-righting probability and decrease energetic cost of dynamic self-righting when its torque and power capabilities decrease (e.g., as battery drains).

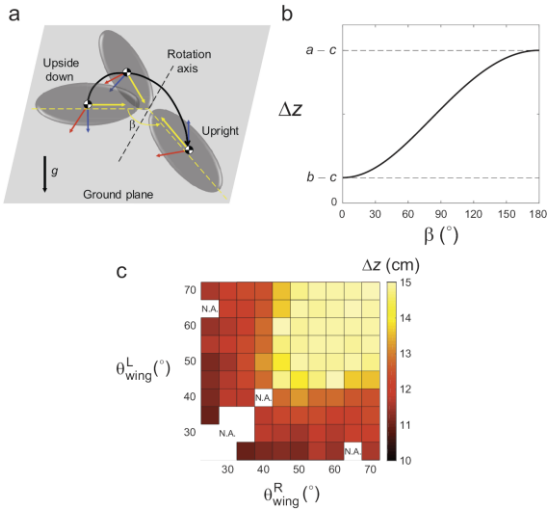


Fig. 14. Geometric model of self-righting using asymmetric wing opening to estimate potential energy barrier height. (a) Schematic of an ellipsoidal body rotating about a fixed axis within the ground plane, resulting in change of heading by β. Yellow, red, and blue arrows represent the fore-aft, lateral, and dorso-ventral directions. (b) Potential energy barrier height, Δz, as a function of β predicted from the model. a, b, and c are the length of the semi axes of the ellipsoid in the fore-aft, lateral, and dorso-ventral directions, respectively. (c) Δz as a function of left and right wing opening magnitudes, θ_{wing}^L and θ_{wing}^R, calculated from the model using the average measured from asymmetric righting experiments.

I. Self-righting in Lizards

We collected ground righting responses from house geckos. In contrast to insect righting, we found that individuals use many kinematically-distinct motions to achieve righting. We hypothesize that these strategies harness one or more of the following three physical mechanisms. Tails and limbs can contact the ground, directly using reaction forces to drive motion (Fig. 15, top sequence). Appendages swung rapidly in the air can induce inertial torques as in aerial righting (tail sweeping, Fig. 15, middle sequence). Bodies can change shape to move the CoM laterally from the contact point, inducing gravitational torques (back-arching, or tail extending, Fig. , bottom sequence). In an individual trial, they may compose several physical mechanisms together (e.g. tail slapping the

substrate halfway through inertial reorientation), increasing the total diversity of response.

We are initially focusing on trials where the tail makes large aerial motions similar to those found in air righting. Geckos using the inertial tail-swing strategy self-right in about 150 ms, approximately the same duration as air-righting geckos. However, two major differences appear – first, the geckos twist their spine significantly in ground righting (up to 80 degrees), while no twisting was reported in air-righting; and second, the tail motions, while still approximately conical, are far larger (up to 500 degrees, vs. ~180 in air-righting).

In planar, zero angular momentum righting, body reorientation is proportional to shape change, suggesting a metric of effectiveness, ξ, defined as the ratio of body to shape rotation [9]. Air-righting geckos rotate their tails through one full sweep of a cone (360 degrees of shape change) to invert the body (180 degrees) for an effectiveness approximately ξ_{air}=0.5. We hypothesized that inertial reorientation in ground righting will be less effective due to parasitic ground reaction forces (deviating from the zero angular momentum approximation). Instead, we found two phases of motion – a low effectiveness region, consistent with our hypothesis of reduced performance due to ground interference, and a high effectiveness region, where the behavior was equivalent to air righting (14, bottom). We will next analyze righting responses where the tail pushes actively against the substrate, hypothesizing much higher effectiveness than in air righting.

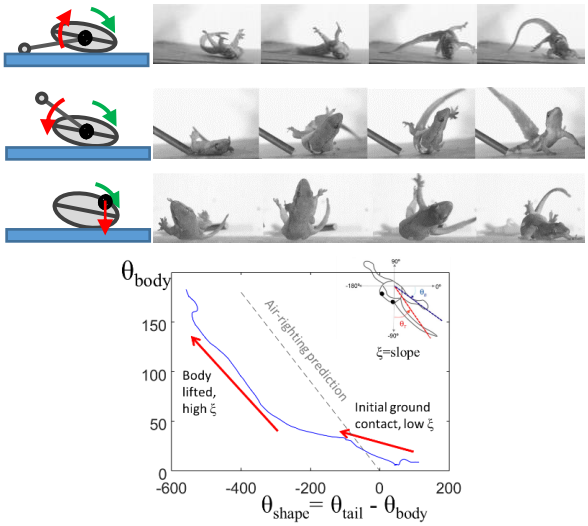


Fig. 15. Lizard self-righting and effectiveness. Top: three distinct strategies for ground righting, corresponding to hypothesized physical explanations. Bottom: Inertial effectiveness, ξ (defined as the ratio of body rotation to shape rotation, slope in this figure) for terrestrial righting.

J. Horizontal Plane Agility in Lizards

We analyzed gait (including limb touchdown and liftoff timing, placement relative to body, and posture) against task-space dynamics in a planar unsteady maneuver. In contrast to previous studies showing small per-step changes in orientation and speed for both animals and robots, we found highly agile (up to 75°/step) maneuvers as lizards turned almost 180° and transitioned to running at near maximal speed. Whereas more gradual maneuvers in animals as diverse as cockroaches [10] and ostriches used forces and gait within the normal range of straight running, we hypothesized that these high-effort maneuvers would run up against the physical constraints (traction, limb force and speed) on locomotor ability and therefore express more dramatic signatures of the force pattern and gait required to accomplish turning.

Defining the task space as average body angle, velocity towards the target, and velocity perpendicular to the target, we evaluated the performance of the observed maneuvers (Fig. 16). Motion through task space was both rapid (up to 1000°/sec, 1-2g acceleration) and consistent, with the average performance capturing around 90% of total variability in task space behavior (measured as Relative Remaining Variance).

To probe the role of gait and the function of individual limbs, we performed perturbation studies where we selectively removed traction beneath one limb at a time (by use of a low-traction surface about ¼ body size). To our surprise, we found no significant effect on performance (time to 150° or 1 m/s), despite foot slips of 3.5cm on average, and additional efforts by animals to avoid the plate.

We found that while the initial (standing) and final (steady trot running) gait were identical in all trials, placement and timing of feet varied greatly during the transition, even in trials with similar task performance. To describe the average gait signal despite observed variance, we built a probabilistic gait diagram by averaging the gait diagrams of individual trials (Fig. 17, top). At each instant in time, the color corresponds to the fraction of trials with a particular limb in contact. In contrast to the stereotyped, physically-constrained performance as we expected, agamas on average reused the steady-running gait (the trot), but with larger duty factor during the transition and high variance in the timing and even number of transition steps, even when controlling for variation in speed and stride frequency.

Variation of limb contact timing was much higher than that of observed task performance. The average gait pattern explained only 53% of variation in the front outer limb, and 68% of variation in the hind outer limb (the most consistent of all limbs). We measured coefficients of variability (CVs) by taking the standard deviation in touchdown timing for a limb relative to its trot partner, divided by the mean step duration for that limb. CVs for the transition stride were 24% and 35% for the inner and outer front limbs, and 22% and 27% for the inner and outer hind limbs, respectively.

While this result seems paradoxical – how could such differing postures at variable times generate forces to produce the same task performance – a closer look at the center of mass dynamics (Fig. 17, bottom) provides an answer. While the average force and torque (blue lines with dashed confidence intervals) are sufficient to produce average task behavior, individual trials (grey lines) use much larger forces that vary about the average – e.g. the "noise" of individual steps is larger than the "signal" actually driving progress in task space. This high variation could be a consequence of changing posture as limbs make and break contact, as well as

the tendency of contacts to provide zero force at touchdown/liftoff and maximum force near mid-step. Regardless, the time-integration of forces into velocity and position acts as a filter, smoothing the high variance of inputs; the result is that many force trajectories ("solutions") achieve approximately the same task behavior.

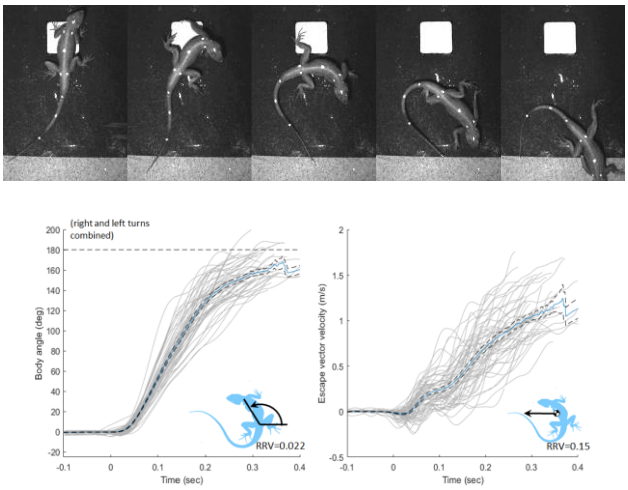


Fig. 16. Lizard turning dynamics. Top: Planar maneuverability task. Bottom: Task space performance (orientation, left, and speed, right) for turn-to-run transition is both rapid and consistent.

It seems extremely unlikely that controller that executes this behavior is strongly "clock-like" as seen in many fast runners, including cockroaches [11] and robots [12]. The stride frequencies observed in these performances (6.3 Hz on average) are likely low enough to permit at least step-to-step feedback modulation of placement and effort. Further, the execution of task space motion does not necessarily dictate the full attention of limb control, even during high-effort maneuvers. The lizard must maintain stability (small fluctuations) in vertical body height, pitch, and roll; the forces used for maneuvering must be net zero impulse in these directions. We hypothesize that the nervous system selects a motor program that maximizes task performance while maintaining stability in "static" modes (e.g. the DOFs not in task space). If the animal could truly operate in task space alone (e.g. confined to a frictionless plane), it could undoubtedly produce higher task performance, with every step aiding the maneuver. In reality, any viable choice may have much higher flexibility than would be expected of a maximal-effort behavior, providing strong robustness to perturbations as a side benefit.

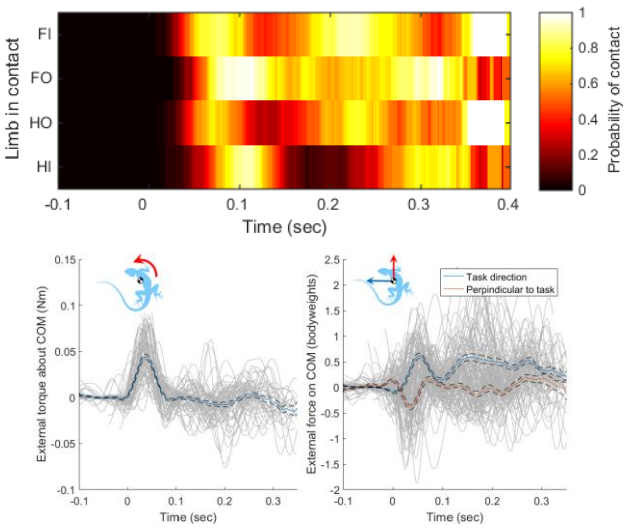


Fig. 17. Lizard turning gait and torques. The limb motions (top; FI, front-inner limb, HO, hind-outer, etc.; black in stance and white in air) that produce the external torques (bottom left) and linear forces (bottom right) are more variable than task space motion.

K. Parallel composition of transitions with sprawled limbs

The most common robot transitions are sequential in nature, composing simple, easily-understood motions in time (e.g. turn first, accelerate second). By contrast, agamas appear to compose their transitions in parallel – they did not separate reorientation and Cartesian motion. COM velocity was initially backwards on average (i.e. towards the prey), and then increased primarily in the lateral axis for the first half of the maneuver. Agamas experienced high lateral COM velocities – almost 40 cm/s on average and frequently over 70 cm/s, over half the observed average task-direction speed. Midway through the turn, this lateral velocity accounted for all the task-direction velocity (as the body axis was exactly perpendicular to the task direction). The angle of misalignment between heading and velocity (i.e. the angle of COM velocity in the body coordinates) remained larger than 90° for the first 100 ms, and exceeded 45° for the first 167ms of the behavior on average; heading and course did not align until nearly the end of the maneuver (over 350 ms from onset of movement). Agamas reached more than 50 cm/s before reaching 90°,i.e. while pointed away from the direction of final motion. Because of the large lateral velocities, agamas were able to avoid generating motion in the direction perpendicular to the task direction – on average, less than 2 cm/s maximum. Turning continued at high rates during this linear acceleration. Heading angular velocity exceeded 600°/s midway through the turn while the lizards simultaneously ran over 50 cm/s on average.

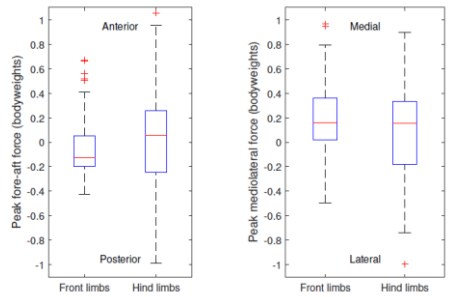


Fig. 18. Peak ground reaction forces in body frame, along the fore-aft axis (left) and mediolateral axis (right). Labels indicate directions of positive and negative forces, e.g. positive mediolateral force is directed towards midline. Extent of box represents the 25th to 75th percentiles and whiskers represent most extreme values not considered outliers.

Agamas' morphology may enable this agile parallel composition. To test the hypothesis that their sprawled posture gives agamas the ability to generate large forces in all directions, we examined the peak forces in the body frame for each limb over all steps and both inside and outside function. Agamas exerted single leg forces anteriorly, posteriorly, and mediolaterally of up to one bodyweight (Fig. 18). Hind limbs generated larger forces than forelimbs in the fore-aft axis, but all limbs generated similar medially-directed forces. Surprisingly, both front and hindlimbs also generated relatively large laterally-directed forces (e.g. the leg pulled the foot towards the midline, rather than pushing outwards) in addition to the medially-directed forces reported in steady running lizards. These pulling forces were largest when the hind limb was inside the turn, where they enabled the largest task-direction linear impulses, well before the animals' body heading aligned with the target direction. The ability of agamas to sustain lateral motion and generate large lateral forces may explain why turning impulse was largely limited to the first half step – once turning is initiated, effort can be directed to linear acceleration much earlier than if they had to complete the turn before accelerating. This strategy may help agamas negotiate the geometric tradeoffs that dominate turning strategy in bipeds, allowing the powerful hind limbs to switch from a turning function to a linear acceleration function as early as their second step.

L. Horizontal Plane Agility in Robots

Our first, two-DOF robot succeeded in producing turns with large per-step rotation. We found stable turning gaits up to ~250 degrees/sec (45 degrees/sec faster than the oscillatory turning gaits found in VelociRoACH and faster unstable gaits. Gaits become unstable when the front limbs touch down early during swing, generating backwards rotational impulses. However, the hind limbs did not appear to aid the turn, and typically hindered motion at least once per step.

Experiments showed that the flexures we used to link action of the hind limbs (also used in VelociRoACH) have poor resistance to lateral leg forces, causing bucking during turning. Our second revision added an additional DOF to the hind limbs, replacing the VelociRoACH-like limbs with RHex-style limbs (Fig. 6). Each limb used a custom-designed servo consisting of a single micro-gearmotor mounted to a rotary potentiometer. Overall mass was not changed relative to the larger motor, geartrain, and flexures of the previous design. Inspired by the lack of clock-like limb control in the turning lizards, we developed a decentralized controller for the new body. Each of the three DOF ran open-loop as long as another limb was in stance phase (estimated by limb position and velocity). If other open-loop limbs were in swing, or would enter swing soon, the lowest-phase limb would be retarded to ensure body support. Once adjacent limbs touched down, the retarded limb would return to open-loop behavior. In contrast to clocked controllers used in many fast-running robots (e.g. RHex), this controller has some similarities to controllers proposed for slow walking (e.g. Walknet [13]), but here is used to generate rapid dynamic motion.

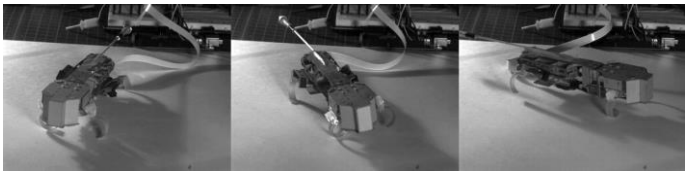


Fig. 19. Trunbot performance. Top: Latest revision of turning robot with independent rear limbs Bottom: Three consecutive touchdowns of front limbs during turning (two strides), with ~ 45 degrees rotation per step.

The new controller produced better turning behavior than the earlier two-DoF version (Fig. 19). Notably, the hind limbs did not appear to hinder movement, and even clearly aided turning on occasion. The new machine turned at speeds up to 450°/sec, and 45°/step, faster than any robot yet reported. We also began preliminary tests on terrain, including discrete steps and unstructured roughness (shredded paper) of about ½ leg length. Turning performance was robust to these obstacles; as long as the front limbs avoided swing-phase contact, turning speeds were unchanged. Front limb swing height is currently the limiting factor on terrain performance; redesigning the linkage to deflect the legs during swing would enable even rougher terrain traversal.

V. Conclusions & Future Work

We have begun to understand the principles of dynamic terrestrial self-righting using wings through systematic experimentation and modeling of a novel winged robot inspired from the discoid cockroach. Analogous to exaptations or co-opting of structures common in biological organisms, our novel winged self-righting mechanism provided rapidly-running small robots with a means to use existing body structures in novel ways to serve new locomotor functions. Inspired by our successful design and discovery of the principles of winged dynamic self-righting, we will continue to develop new robot prototypes that will have integrated ability to both traverse obstacles and self-right if flipped over. We will continue to explore the more multifarious solutions used by lizards in self-righting by changing tail affordance (removing or altering substrate) and through the comparative method by examining other species.

Our study of planar maneuverability revealed surprising features of unsteady locomotion, countering our expectations

that high-effort maneuvers should be stereotyped in gait and force pattern. Our results hint at a control architecture yet undescribed. Inspired by the turning performance of the lizards, we built a turning robot similar in size and design to other legged MAST platforms. Equipped with a decentralized controller, this robot achieved the fastest legged turning ever described. We will next return to the active spine design and attempt to improve on this performance by adding coordinated core and tail actuation.

VI. References

[1] K. H. Low, T. Hu, S. Mohammed, J. Tangorra, and M. Kovac, "Perspectives on biologically inspired hybrid and multi-modal locomotion," *Bioinspir. Biomim.*, vol. 10, no. 2, p. 20301, 2015.

[2] C. Li, A. O. Pullin, D. W. Haldane, H. K. Lam, R. S. Fearing, and R. J. Full, "Terradynamically streamlined shapes in animals and robots enhance traversability through densely cluttered terrain," *Bioinspir. Biomim.*, vol. 10, no. 4, p. 46003, 2015.

[3] C. C. Kessens, D. C. Smith, and P. R. Osteen, "A Framework for Autonomous Self-Righting of a Generic Robot on Sloped Planar Surfaces," *IEEE Int. Conf. Robot. Autom.*, pp. 4724–4729, 2012.

[4] C. Li, T. Wohrl, H. K. Lam, and R. J. Full, "Body and appendage morphology affect ground-based self-righting of cockroaches," *J. Exp. Biol.*

[5] S. Jay, E. S. Vrba, S. Paleobiology, N. Winter, S. J. Gould, and E. S. Vrba, "Exaptation-A Missing Term in the Science of Form," *Palaeontology*, vol. 8, no. 1, pp. 4–15, 1982.

[6] A. Jusufi, D. I. Goldman, S. Revzen, and R. J. Full, "Active tails enhance arboreal acrobatics in geckos.," *Proc. Natl. Acad. Sci.*, vol. 105, no. 11, pp. 4215–9, Mar. 2008.

[7] C. Li, C. C. Kessens, R. S. Fearing, and R. J. Full, "Mechanical principles of dynamic terrestrial self-righting using wings," *Adv. Robot.*, 2017.

[8] G. Domokos and P. L. Várkonyi, "Geometry and self-righting of turtles.," *Proc. R. Soc. London B Biol. Sci.*, vol. 275, no. 1630, pp. 11–7, 2008.

[9] T. Libby, A. M. Johnson, E. Chang-Siu, R. J. Full, and D. E. Koditschek, "Comparative Design, Scaling, and Control of Appendages for Inertial Reorientation," *IEEE Trans. Robot.*, vol. PP, no. 99, pp. 1380–1398, 2016.

[10] D. Jindrich and R. J. Full, "Many-legged maneuverability: dynamics of turning in hexapods," *J. Exp. Biol.*, vol. 202, pp. 1603–1623, Jun. 1999.

[11] S. Sponberg and R. J. Full, "Neuromechanical response of musculo-skeletal structures in cockroaches during rapid running on rough terrain.," *J. Exp. Biol.*, vol. 211, no. Pt 3, pp. 433–446, Mar. 2008.

[12] U. Saranli, M. Buehler, and D. E. Koditschek, "RHex: A Simple and Highly Mobile Hexapod Robot," *Int. J. Rob. Res.*, vol. 20, no. 7, pp. 616–631, 2001.

[13] M. Schilling, T. Hoinville, J. Schmitz, and H. Cruse, "Walknet, a bio-inspired controller for hexapod walking," *Biol. Cybern.*, 2013.

Real-time Force Feedback in Heterogeneous Terrain

Task MCE 2.4

Carlos S. Casarez
Dept. of Mechanical Engineering
University of California at Berkeley
Berkeley, CA
casarezc@berkeley.edu

Can Koc, Cem Koc, Brian Su, and Ronald S. Fearing
Dept. of Electrical Engineering and Computer Sciences
University of California at Berkeley
Berkeley, CA
ronf@eecs.berkeley.edu

***Abstract*—As palm-sized legged robots enter the natural world, they will encounter challenging terrain with obstacles of varying size scales. A low-cost 6-axis force/torque sensing shell has been developed to measure the interactions between a terrestrial robot and densely cluttered environments. Experiments testing the tactile shell integrated into the VelociRoACH robot platform show that the novel sensor can characterize the drag energy required to pass through narrow channels, which affects the specific resistance of locomotion. These experiments lay the groundwork for planning efficient paths through heterogeneous terrain. In environments with increased terrain complexity, VelociRoACH can encounter steep drop-offs and rough terrain that may flip the robot over, causing it to become immobile. To address this potential failure mode, an active tail for self-righting is added to VelociRoACH. Physical analysis of terrestrial self-righting results in design requirements and righting time predictions for a shelled robot with an active tail. Open-loop experiments on terrain with varying friction and roughness show that VelociRoACH can self-right using tail contact in as little time as 256 ms. Finally an autonomous self-righting experiment on an obstacle with multiple step drops demonstrates that the robot can detect inversion and rapidly self-right while walking on challenging terrain.**

***Index Terms*— Tactile sensing, legged robots, active tail**

I. Introduction

A. Background and Technical Challenges

MAST ground platforms will be faced with challenging terrain outside the laboratory. As shown in Fig. 1a, this heterogeneous terrain will not only interact with legs and feet, but will present significant contact forces on the robot body. These contact forces can not only resist robot motion, but could also destabilize the robot as shown in Fig. 1b.

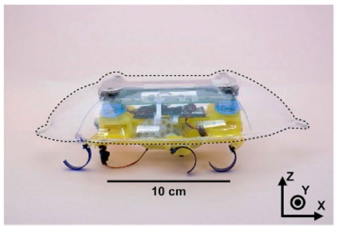


(a)

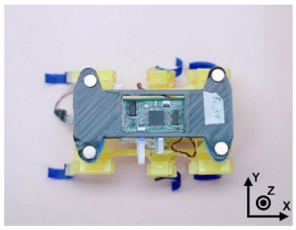


(b)

Fig. 1. a) Example challenging terrains with high resistance to robot traversal. b) Example traversal using passive reorientation through shell interactions with terrain [1].



(a)



(b)

Fig. 2. VelociRoACH robot with (a) removable force/torque sensing shell attached, (b) force-sensing shell removed showing magnetic fixtures and novel planar low-cost optical 6-axis force/torque sensor.

Through the use of a force sensing shell, a robot can be controlled to better negotiate this type of challenging terrain. The specific technical challenges are providing a robot with a force sensing shell, estimating body contact forces in real time, and stably controlling robot thrust and orientation.

There has been a wide variety of work using contact sensors to navigate environments, particularly using mechanoreceptors such as antennae or whiskers. A cockroach-inspired wall-following control system was developed for a small wheeled robot by Cowan et al. [2] and J. Lee et al. [3]. In this situation, the robot is basically in free space, and does not experience any resistance from interacting with the environment. In effect, the antenna can act as a proximity or position sensor (albeit with interesting internal sensor dynamics). Similarly, an array of binary contact switches can estimate the shape of terrain on the bottom of a robot, as demonstrated by Karras et al. [4]. In contrast, traversing environments such as those shown in Fig. 1 requires force sensing, as the robot body will be in contact with the environment, perhaps at multiple points.

Most robot path planning work is in avoiding contact with obstacles, or minimizing contact through the use of position or proximity sensors. With a cluttered environment, the robot will be making contact continually. While many small mobile robots have protective shells, the only prior work found in a force sensing shell is the work by Tsuji et al. [5], [6]. In this work, the sensor function was demonstrated, but the force data was not used in any contact control application. Unlike a distributed tactile sensor array (e.g. Fearing and Binford [7]), the contact force sensor approach provides effectively the centroid of the contact location (Bicchi et al. [8]). There is a large body of work with robot arms in contact, such as the operational space force control system pioneered by Khatib [9].

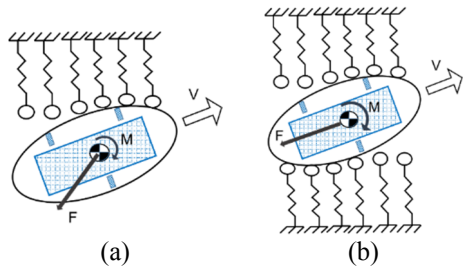


Fig. 3. Net force sensing in an environment with distributed contact. The robot is walking in the direction *V*, with 3 legs in ground contact. In (a), the robot detects a drag force resisting its forward velocity and a pushing force away from the wall. In (b), the drag force is detectable, but the symmetric squeezing force is not detectable.

Commercial force/torque sensors, such as the ATI Nano17, could be used in larger robots, but with interface electronics they typically have mass of several hundred grams and cost over $5000, which is incompatible with small low-cost MAST platforms. A low-cost, low resolution (order 1%) force/torque sensor has been designed using a novel planar optical sensor array as shown in Fig. 2. A prototype of this sensor has been integrated with the VelociRoACH mechanical, electrical, and software system, as presented by Goldberg and Fearing [10]. The sensor consists of a planar printed circuit board with optical proximity sensors. As the sensor prototype is already functioning, the main effort through this MAST task is to develop legged locomotion control considering body forces. As described below, body forces can increase drag, or in the worst case, cause the robot to tip over.

There are several fundamental differences between distributed contact sensing and net force/torque sensing, as is shown in Fig. 3. The net force/torque sensing approach cannot resolve any set of forces which sum to zero, for example a squeezing force (Fig. 3b). On the other hand, a distributed contact sensor could recover the shape of the contact in Fig. 3a, but it would generally be difficult to recover contact stresses without performing a deconvolution operation, such as in Fearing and Binford [7]. One strong advantage of the force/torque sensor is that it provides a direct measurement of lift/drag forces, and pitch/roll moments which are important for body stabilization in complicated media. Ideally, a robot would combine both distributed and net contact force sensors, as each sensing modality has cases where it is more informative.

With the addition of a force-sensing shell, underactuated legged robots can expand their application realm to cluttered, unstructured terrain. Consequently, these robots will encounter steep drop-offs and may flip over when moving through dense clutter. Therefore, the capability to self-right is essential for maintaining leg contact with terrain. Previous robotic systems have demonstrated the use of low-DOF appendages for self-righting both in the air and on the ground. Libby, Chang-Siu, and Johnson et al. analyzed and tested a lizard-inspired strategy for aerial self-righting, in which a robot pushes on the inertia of a massive 1 or 2 DOF tail to reorient its body [11], [12], [13]. Krummel et al. designed a robot inspired by a horseshoe crab, which can traverse surf zones and self-right by pushing a tail appendage against the ground [14]. Li et al. demonstrated

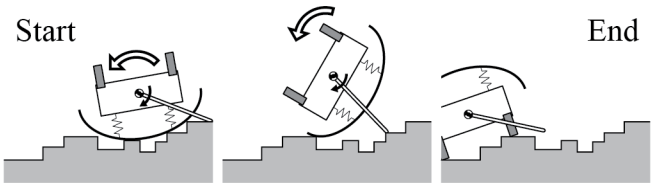


Fig. 4. Tail-assisted self-righting of a legged robot with tactile shell on rough terrain.

cockroach-inspired terrestrial self-righting through the outward unfolding of two servomotor-driven wings [15]. Kessens et al. developed a contact planning framework for performing self-righting maneuvers with low-DOF appendages [16], [17]. The principles of minimally-actuated self-righting and previous successful robot designs will be considered in order to integrate an active tail for self-righting into the design a VelociRoACH robot with tactile force/torque sensing shell.

B. Research Objectives

The objective of this research task is to develop a closed-loop force control for VelociRoACH [18] to enable both path following and body reorientation control, such as needed in righting. There are a variety of ambulation situations where mobile robot force control will be advantageous. Specifically:

1. Controlling contact with terrain, for example avoiding the problems of high-centering, or getting stuck by a low overhang. As the shell is connected to the chassis through the force/torque sensor, contact forces on the bottom surface of the robot can be monitored, and an alternate path could be found if the robot is in danger of getting stuck.
2. Terrain which has a similar visual appearance could have variable resistance. A force control strategy can be developed to negotiate paths through terrain. This strategy could also work in low-light conditions.
3. A particularly challenging situation where terrain contact is critical is righting control for an inverted robot, as shown in Fig. 4. On a flat surface, righting strategies using an extra tail-like appendage can be pre-calculated exactly (e.g. Kessens et al. [16], [17]). However, in very rough terrain, a strategy which works on flat terrain may not generate sufficient moment for righting the robot. In these cases, the righting appendage can be controlled more effectively (e.g. avoiding jamming of the robot) by considering the net contact force on the robot shell.

C. Technical Approach Summary

In this report, two separate sub-problems of legged robot interaction with rough unstructured terrain are addressed.

The first sub-problem is experimentally identifying the energetic cost of traversing narrow channels with grass-like beams that exert drag forces on a legged robot. Characterizing the resistance to motion of different cluttered terrain types will provide measurements guiding the development of a control strategy for navigating heterogeneous terrain.

The second sub-problem is achieving terrestrial self-righting of a shelled robot with an active tail. As part of this sub-problem,

quasi-static force analysis determines the motor torque required for active self-righting, dynamic simulation predicts open-loop self-righting time, and a legged robot with active tail performs open-loop self-righting on various terrain types, including centimeter-scale rocks. In addition, the robot autonomously recovers from inversion on terrain with multiple step drops.

II. Related Work

A. Synergy with Existing Works and External Collaborations

This task has mutually benefitted from interactions with Task MCE 2.1: Robophysical Approach to Discovery of Principles of Effective Interaction of Ambulatory Robots in Multicomponent Environments, and Task MCE 2.2: Bio-inspired Principles of Extended Terrain Mobility.

The self-righting work of Task MCE 2.3: Bio-inspired Principles of Transitions was directly relevant to this task, as measure and control of body contact forces is necessary for reliable "self-manipulation" on complicated terrain. The ability to measure body contact forces and locations can also provide insights into details of the collisions and contacts driving orientation transitions. A possible collaboration on self-righting with sensing will be explored with C. Kessens at ARL.

B. Technology Transitions Summary

There are many potential opportunities for transitioning the basic research of this task to mobile robot technology. Tactile sensing of the environment can provide valuable force feedback to aid MAST platforms in the navigation of cluttered terrain. Furthermore, locomotion studies characterizing the drag energy and specific resistance of passing through cluttered terrain can inform the design of streamlined robot chassis or added shells that enable MAST platforms to move through challenging natural terrain with more energetic efficiency.

Active tails can expand the capabilities of MAST platforms beyond controlled laboratory settings. As is shown in this report, an active tail can recover VelociRoACH from an inverted posture that would have previously been an irrecoverable failure state. With the addition of a single motor and a lightweight tail, MAST platforms with similar limitations as VelociRoACH could more readily be deployed in environments with steep drop-offs and terrain of varying roughness scales.

III. Approach

A. Tactile Shell Terrain Characterization Experiments

For experiments moving through cluttered terrain, a VelociRoACH with tactile shell (shown in Fig. 2) can characterize the resistance to motion. The robot traverses terrain and interacts with the environment while the force-sensing shell records forces in the x, y, and z directions. Experiments were performed with the robot following an alternating tripod gait at a stride frequency of 1 Hz on low-pile carpet.

Inspired by the model of cluttered terrain from Li et al. [1] compliant plastic beams are placed along two parallel walls with a channel length $l_c = 25.4$ cm. By varying the channel width w_c, the drag force exerted on the robot can be adjusted. The tips of the grass-like beams are spaced apart laterally to form channel

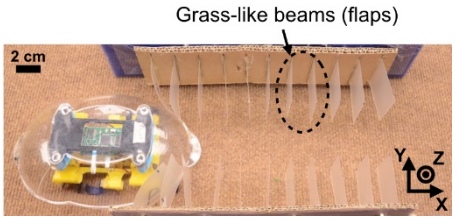


Fig. 5. Experimental setup for VelociRoACH with tactile shell traversing a high-density channel with flaps spaced at a width of 6 cm.

widths from 4 cm to 10 cm, with an interval of 2 cm. Since the robot is 10 cm wide at its widest point, these channel widths correspond to maximum lateral beam deflections (in the y direction) of 0 cm to 3 cm, with an interval of 1 cm. As a control, a "free" channel with no flaps was also tested.

The experimental setup for terrain drag characterization is shown in Fig. 5. The robot approaches the channel, enters the high-density region, and eventually exits the channel. Telemetry data is recorded from the force-sensing shell to compute the forces acting on the shell in the x, y, and z directions (F_x, F_y, F_z). In addition, the applied voltage and back-emf voltage of the drive motors of the robot are used to compute the power draw of each motor over the course of the trial. From the telemetry data, the drag energy E_{drag} resisting motion can be computed:

$$E_{drag} = \bar{F}_x l_c \,, \tag{1}$$

where $\bar{F}_x$ is the average drag force resisting the robot's forward progress over the duration of the trial. Eq. 1 gives a measure of the additional work that the robot has to perform to traverse the channel compared to open flat terrain. In order to compare the performance of the robot across different experimental trials, the specific resistance η is computed:

$$\eta = \frac{\bar{P}}{mg\bar{v}} \,, \tag{2}$$

where $\bar{P}$ is the time averaged power consumption of both drive motors and $\bar{v}$ is the time averaged forward velocity of the robot over the duration of the trial. The specific resistance in Eq. 2 is a dimensionless quantity that is a ratio of the electrical power input to the task-based power output, which is the robot weight mg multiplied by the average forward velocity.

B. Tail-assisted Self-righting Analysis

An end goal for tail-assisted self-righting of a VelociRoACH robot with tactile shell is to be able to autonomously detect inversion conditions and then perform a controlled self-righting maneuver with an active tail to flip the robot back onto its legs, as shown in Fig. 4. Eventually, self-righting will occur on sloped and unstructured terrain, with the tactile shell giving important feedback on whether the tail is making progress towards self-righting based on the measured contact forces of the shell against the terrain. A dynamic model of tail-assisted self-righting on flat terrain with uniform friction properties is developed to determine tail design requirements and to predict the time required to self-right with an open-loop tail swing.

For the dynamic model of tailed self-righting, Fig. 6 shows the mass and geometry parameters of a VelociRoACH robot with tactile shell and tail. The robot body with shell has mass m_b

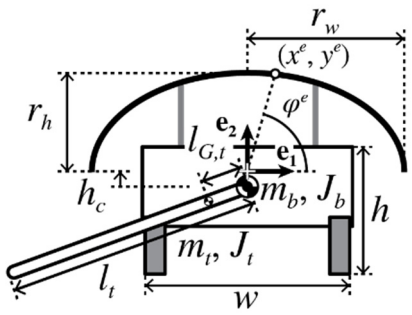


Fig. 6. Back view diagram showing mass and geometry parameters of a VelociRoACH robot with shell and active tail.

TABLE I. VELOCIROACH WITH SHELL AND TAIL PARAMETERS

Parameter name	Symbol	Value
Body mass	m_b	76.7 g
Body mass moment of inertia	J_b	583 g-cm^2
Overall width, height	(w, h)	(7.5, 4.5) cm
Shell major, minor axis radius	(r_w, r_h)	(5.5, 3.5) cm
Shell center offset height	h_c	0 cm
Tail mass	m_t	0.8 g
Tail mass moment of inertia	J_t	135 g-cm^2
Tail length, C.o.M. distance	$(l_t, l_{G,t})$	(9, 1.5) cm
Tail stall torque	τ_s	29.5 mN-m
Tail no-load speed	ω_{nl}	77 rad/s

and mass moment of inertia about its roll axis J_b. The robot without shell has overall width w and overall height h. The robot's rigidly mounted shell is approximated as half an ellipse in the roll plane of the robot, with major axis radius along the width of the robot r_w, minor axis radius along the height of the robot r_h, and center offset height above the center of mass (C.o.M.) of the robot h_c.

A rigid tail is driven by a motor in the roll plane of the robot, with an assumed pivot point about the C.o.M. for simplicity of analysis. The tail has mass m_t and mass moment of inertia about its C.o.M. J_t. The tail C.o.M. is located a distance $l_{G,t}$ from the tail pivot. Model parameters for this robot are given in Table I.

Fig. 7 shows free body diagrams of tail-assisted self-righting. The analysis of self-righting is performed in the roll plane of the robot with world-fixed Cartesian frame $\{\mathbf{E_1}, \mathbf{E_2}, \mathbf{E_3}\}$. The forces that act on the two body system include gravitational forces ($\mathbf{g} = -g\mathbf{E_2}$, $g = 9.81$ m/s^2), body contact force $\mathbf{F_b}$, and tail contact force $\mathbf{F_t}$. In Fig. 7a, $m = m_b + m_t$ is the total system mass. The net inertial force due to the linear acceleration of the system C.o.M. is $-M\mathbf{a_G} = -m_b\mathbf{a_{G,b}} - m_t\mathbf{a_{G,t}}$.

For a given body roll angle θ_b relative to $\mathbf{E_1}$, the kinematics of a rolling ellipse provide constraints on the velocity and acceleration of the robot body. The tail angle θ_t and its time derivatives are also constrained so that the tail maintains contact with the ground. The input motor torque τ in Fig. 7b is limited by a DC motor model in which torque magnitude linearly decreases with increasing relative tail rotation speed $|\dot{\theta}_t - \dot{\theta}_b|$. The contact forces $\mathbf{F_b}$ and $\mathbf{F_t}$ have directions that are limited by independent friction coefficients μ_b and μ_t, where each friction cone angle α satisfies $\mu = \tan\alpha$. If a contact is sliding, then the respective contact force is at the edge of the friction cone. If a contact is sticking, then the respective contact force is within the friction cone. At least one contact must be sliding in order for the dynamics equations to produce a feasible solution. For full details on the formulation of the system dynamics, refer to [19].

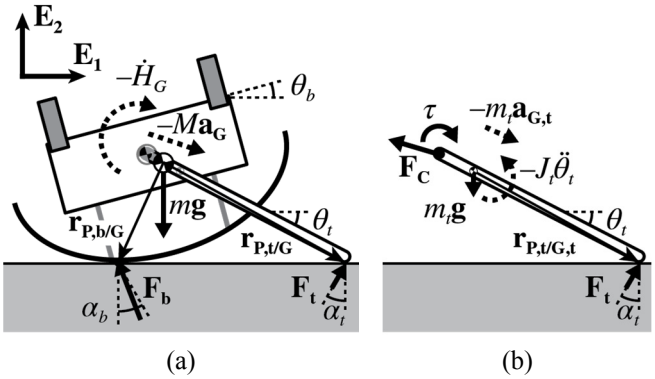


Fig. 7. Roll plane diagrams of the forces and moments that act on a tailed robot during a dynamic self-righting maneuver. (a) Free body diagram of the overall system and (b) free body diagram of the tail.

In the results section, a quasi-static analysis gives minimum torque requirements for the tail motor and the dynamic system is simulated to predict self-righting time with a contacting tail.

C. Tailed Robot with Self-righting Capabilities

The legged robot in Fig. 8 was developed for tail-assisted self-righting experiments. The base platform is a modified VelociRoACH robot [18] with a mass of 53.6 g. The robot is powered by a single cell 300 mAh LiPo battery (*Lectron Pro*). The kinematic linkages of the robot have PET plastic as the structural material and ripstop nylon as the flexural material. To drive the robot legs with added payload, the leg transmissions have four times the torque of the original VelociRoACH, at the expense of a four times reduction in stride frequency. The robot has a 0.254 mm thick polycarbonate shell (adds 10.6 g) and a tail motor (*Pololu* 30:1 reduction high power micro metal gearmotor) that drives an interchangeable tail (adds 13.3 g). The total mass of the robot with added shell and tail is 77.5 g. The tail is a flat carbon fiber beam (0.8 g with attachment hardware) with a width of 3 mm and an effective length of 9 cm.

IV. RESULTS

A. Resistance to Motion of Cluttered Terrain

Telemetry data from a representative trial of VelociRoACH with a tactile shell traversing a channel of compliant beams with a lateral deflection of 3 cm is shown in Fig. 9. The shaded red

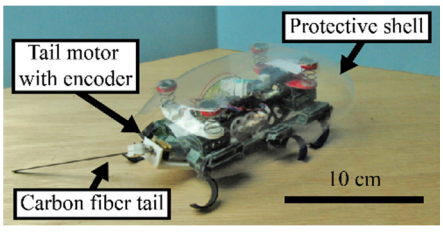


Fig. 8. Annotated image of VelociRoACH with protective shell and tail.

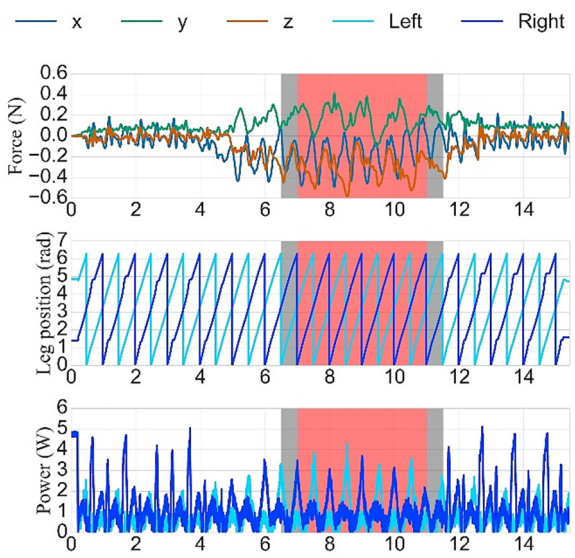


Fig. 9. Telemetry data from VelociRoACH traversing a cluttered channel with a lateral beam deflection of 3 cm. From top to bottom: forces in the x, y, and z directions recorded by the force-sensing shell, left and right leg positions, and left and right motor power from robot telemetry. The shaded red region marks when the robot is entirely within the channel.

region shows the time when the robot is in the channel. The time trajectories of forces (top plot in Fig. 9) show that the amplitude of force variations in the x, y, and z directions increases upon entering the channel. As expected, the forces exerted on the shell vary with a fundamental period of either 0.5 s or 1 s (half stride period or full stride period), which is synchronized with each leg set's tripod step at 1 Hz phased 180 degrees apart (middle plot in Fig. 9). There is an average increase in the measured drag force along the channel (negative x direction) as well as an increase in the downward force pushing on the robot (negative z direction). However, because this distributed contact case matches Fig. 3b, there is no significant increase in the average lateral force perpendicular to the channel (y direction). Because the tactile sensor is connected to the body of the robot by a spring, small oscillations from the robot's movement are captured by the force/torque sensor and result in some noise, as seen in the force plot before and after the shaded region.

Now, the drag (x direction) and lift (z direction) forces are compared across representative trials of the five tested beam deflection cases in Fig. 10. As expected, the free case (Fig. 10a) exhibits a uniform oscillatory response, with small variations in force due to movement of the spring-mounted sensor coupled with the robot's legged movement. For the cases with beam deflections from 0 cm to 3 cm (Figs. 10b–10e), a fundamental oscillation period of 0.5 s is maintained in the drag responses, with the average drag force inside the channel (shaded red region) becoming more negative as the beam deflection increases. The deflected beams also exert negative lift (downward z force) on the robot as it passes through the channel, although with less pronounced oscillations at a fundamental oscillation period of 1 s. For the large beam deflection cases of 2 cm (Fig. 10d) and 3 cm (Fig. 10e), the negative lift steadily decreases to exceed the magnitude of the drag force as the robot approaches the middle of the channel, and then increases toward

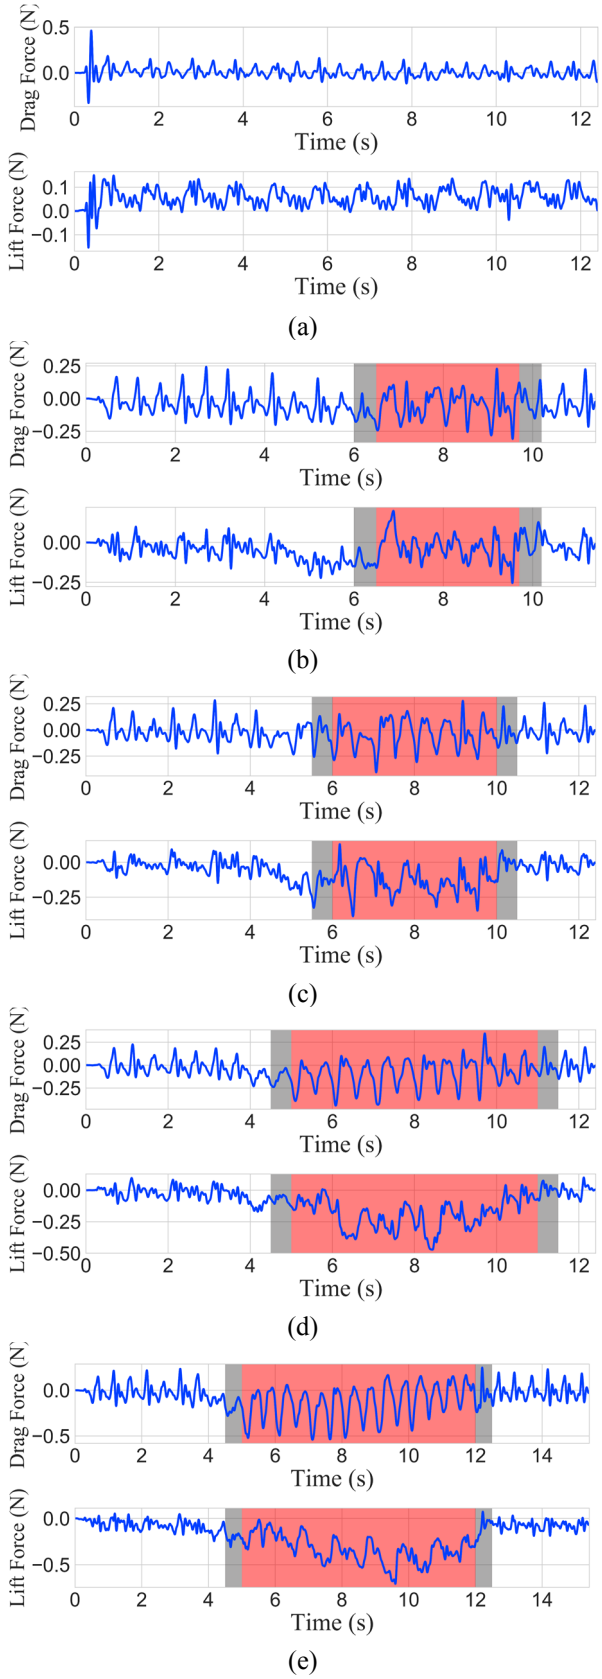


Fig. 10. Drag (x direction) and lift (z direction) forces during representative trials of VelociRoACH traversing a channel with (a) free, (b) 0 cm, (c) 1 cm, (d) 2 cm, and (e) 3 cm lateral beam deflection cases. The shaded red region marks when the robot is entirely within the channel.

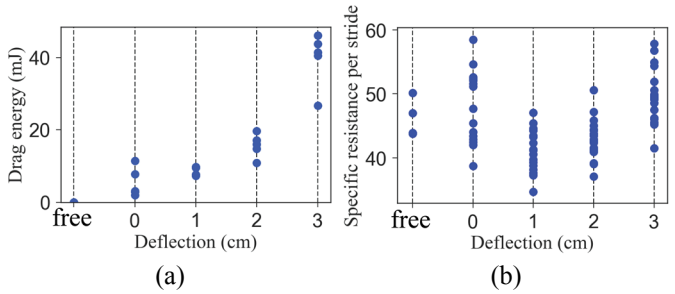


Fig. 11. (a) Average drag energy over the course of each trial and (b) specific resistance per leg stride in aggregated trials for VelociRoACH traversing high-density channels with varying lateral deflection of grass-like beams.

zero as the robot leaves the channel. These plots also show that as beam deflection increases, the increased magnitude of the body forces exerted on the robot cause it to take a longer time to traverse the channel.

Using 5 trials for each beam deflection case, data was collected as in Fig. 9. Eq. 1 was used to compute drag energy and Eq. 2 was used to compute specific resistance. Note that the specific resistance was computed as a "per stride" metric, with input power averaged over each individual leg stride and velocity averaged over the entire trial. Fig. 11a shows the drag energy measured by the robot for channels with varying beam deflection, including the free case with no flaps. The plot shows a nonlinear increase in drag energy as beam deflection increases.

In Fig. 11b, the specific resistance per leg stride is shown for varying beam deflection. Each data point is the specific resistance for a leg stride extracted from 1 of 5 trials for each beam deflection case. The robot incurs greater drag energies at larger beam deflections and therefore performs more leg cycles in large deflection experiments. As a result, there are 12, 16, 18, 19, and 22 points for the free, 0 cm, 1 cm, 2 cm, and 3 cm experiments respectively. Fig. 11b shows an overall increase in specific resistance as beam deflection increases. In a deviation from this overall trend, the 0 cm deflection data points show increased spread and an average specific resistance of 48.0, which is closer to the average specific resistance of the 3 cm experiments, 49.5, than to the average specific resistance values of 45.2, 41.1, and 42.9 for the free, 1 cm, and 2 cm experiments. The increased spread in the specific resistance per leg stride for the 0 cm deflection case can be explained by the robot's tendency to veer towards one of the flap walls and subsequently collide with the wall.

Moreover, the fact that the average specific resistance does not strictly decrease with decreasing beam deflection suggests that squeezing through a channel with some downward and backward force can result in less energetic cost than running in the free condition. The average specific resistance is lower in the 1 cm beam deflection experiment than in the free experiment. However, if the channel is too narrow, such as in the 3 cm beam deflection experiment, then running through a channel becomes more energetically costly than running outside a channel. This result suggests that robots may be able to leverage environment contact to navigate with more energetic efficiency.

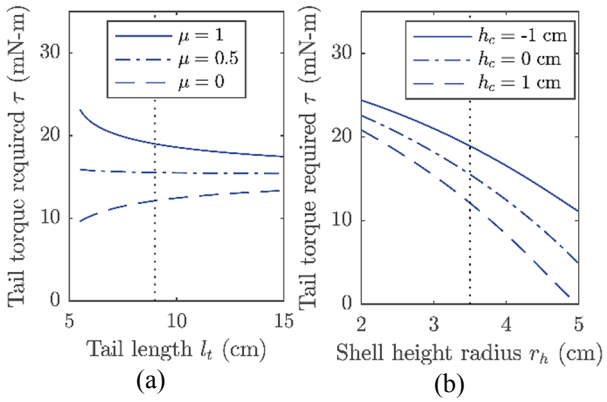


Fig. 12. Tail torque required for quasi-static self-righting of a robot with an elliptical shell on flat terrain. (a) Varying tail length and tail-ground sliding coefficient of friction, with robot shell parameters held constant at $r_h = 3.5$ cm and $h_c = 0$. (b) Varying shell height radius and offset height above the C.o.M., with $l_t = 9$ cm and $\mu = 0.5$.

B. Self-righting Tail Design and Motor Selection

In order to select a tail length and tail motor for self-righting, a quasi-static analysis was conducted on the system in Fig. 7. With a quasi-static assumption, the tail motor maintains static equilibrium throughout the righting process. In quasi-static self-righting, the tail rotates the body from an inverted configuration ($\theta_b = 0°$) to a critical body angle ($\theta_b = \theta_b^*$). Past this critical body angle, gravity tends to tip the robot onto its legs. The maximum tail motor torque over the range of admissible body angles provides a quasi-static limit on the stall torque τ_s required for tail-assisted self-righting on a horizontal surface.

First, the quasi-static tail torque required for self-righting is determined for varying tail length l_t and varying coefficient of friction $\mu = \mu_b = \mu_t$, with remaining parameters fixed at the values in Table I. These results are plotted in Fig. 12a. To reduce the tail torque required to self-right on high-friction surfaces, a tail with $l_t = 9$ cm was selected for VelociRoACH.

Next, the tail and friction parameters of the quasi-static model are held constant while the shell minor axis radius r_h and the shell offset height above the C.o.M. of the robot h_c are varied. The shape variation results in Fig. 12b suggest that the shell should extend as high as possible above the C.o.M. to reduce the torque required for self-righting. However, increasing the overall height of the robot can prevent the robot from entering low clearances. This size constraint limits the shell height parameters of VelociRoACH to $r_h = 3.5$ cm and $h_c = 0$.

Based on the quasi-static analysis, a 30:1 reduction gearmotor (*Pololu*) was selected to drive the tail, which has an output stall torque of $\tau_s = 29.5$ mN-m and a no-load speed of $\omega_{nl} = 77$ rad/s at a voltage of 3.7 V. This motor provides enough torque for the robot to quasi-statically self-right on flat terrain over the range of considered friction conditions.

C. Dynamic Tail-Assisted Self-righting Time Predictions

The system of second-order differential equations governing tail-assisted self-righting dynamics was numerically simulated using `ode45` in MATLAB, with events handling transitions between sliding and sticking contacts as well as detecting loss of

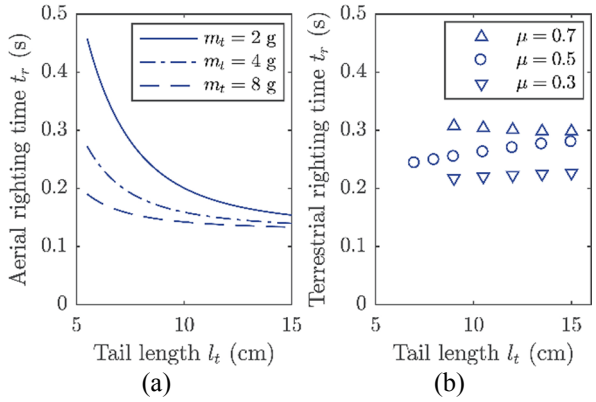


Fig. 13. Comparison of the time to perform a 180° roll reorientation of a robot body for (a) aerial self-righting with an inertial tail of varying mass and length and (b) terrestrial righting with a low-mass tail of varying length on surfaces with varying coefficient of friction $\mu = \mu_b = \mu_t$.

contact when normal forces reach zero. Body angle, tail angle, and tail torque trajectories for the simulation of tailed righting on a wood surface are shown as a dashed line in Fig. 15.

The dynamic simulation's prediction of terrestrial righting time is now compared to aerial righting. Fig. 13a shows the time t_r to complete aerial righting determined from the model in [13] for varying tail length l_t and tail end mass m_t. For aerial righting, t_r decreases as l_t and m_t increase (which increases the inertial leverage of the tail), with greater effects on t_r for tails with l_t ranging from 5 cm to 10 cm.

Assuming uniform friction ($\mu = \mu_b = \mu_t$), the dynamic simulation was used to predict the terrestrial righting time for varying tail length l_t and coefficient of friction μ. In contrast to aerial righting, the terrestrial righting time results in Fig. 13b show that t_r does not change much with varying l_t. The greatest effect of tail length is observed in the $\mu = 0.5$ case, in which increasing l_t from 7 cm to 15 cm increases t_r by 37 ms. Righting time is far more sensitive to the coefficient of friction μ—for a tail length $l_t = 9$ cm, increasing μ from 0.3 to 0.7 increases t_r by 90 ms. Overall, terrestrial righting over the range of l_t and μ values is 84–174 ms slower than the fastest aerial righting case.

D. Open-loop Tail-assisted Self-righting Performance

The self-righting performance of VelociRoACH was tested by executing an open-loop tail control sequence while inverted. A sequence of high-speed video frames are shown for a representative open-loop self-righting trial in Fig. 14. In this experiment, the robot started inverted on a flat wood surface with its tail against the ground. Using a PID controller with feedback on tail position (measured with a magnetic rotary encoder), the robot applied an input voltage to swing its tail 180° clockwise. As pictured in Fig. 14, the robot pushes its tail against the ground, rolls counterclockwise, and self-rights. Telemetry data for ten repeated open-loop self-righting trials on wood are shown in Fig. 15. The telemetry plots show the time trajectories of the body roll angle (measured with the onboard accelerometer and gyroscope), the tail rotation angle relative to the body (the zero reference for θ_t is the initial contact angle), and the tail torque. The tail torque is calculated using a linear DC motor model and the tail encoder measurements.

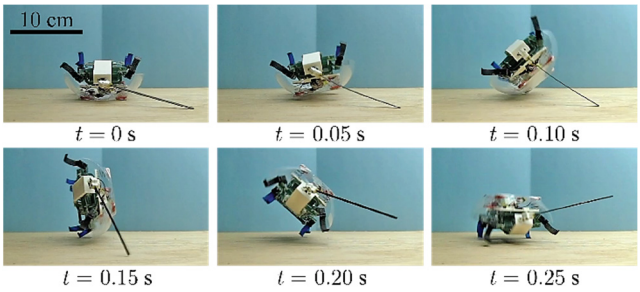


Fig. 14. High-speed camera frame sequence of a VelociRoACH robot with shell and active tail performing a dynamic self-righting maneuver on wood.

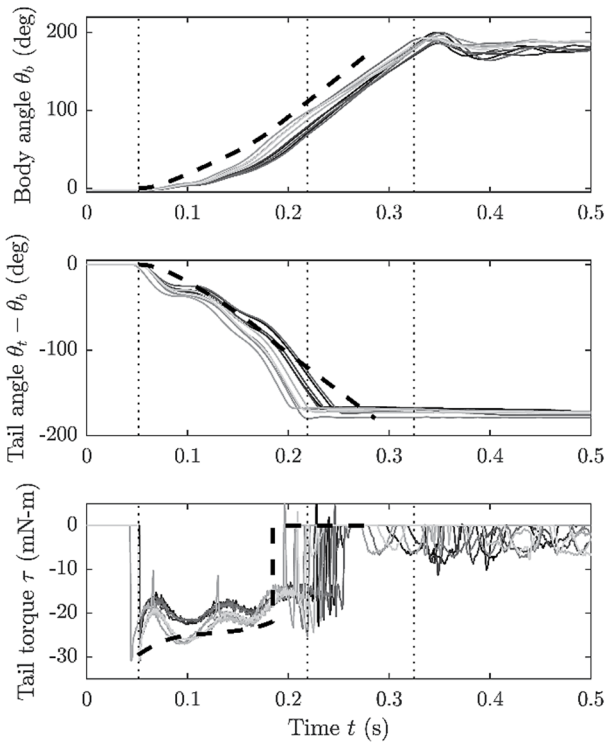


Fig. 15. Robot telemetry for ten repeated trials of open-loop tail-assisted self-righting on wood. The vertical dotted black lines mark the start of righting, the end of the period of tail work, and the end of righting. The dashed black line is the dynamic simulation result.

For the wood surface experiments, the average righting time was 274 ms, with a standard deviation of 7.71 ms. Compared to the dashed simulation result ($\mu_b = 0.4$, $\mu_t = 0.5$), open-loop self-righting on wood was on average 39 ms slower. The planar model may underpredict righting time because in the experiments, the tail produced pitch and yaw motions outside of the roll plane. These higher dimensional dynamics could contribute to the more gradual roll angle acceleration of the robot in the experiments relative to the model.

Open-loop tail-assisted self-righting was tested on the seven surfaces shown in Fig. 16. Table II shows experiment righting times (average ± standard deviation) and success rates over ten trials for self-righting on each surface. The body and tail surface friction coefficients were determined from inclined plane tests measuring static friction—body friction μ_b was measured for the polycarbonate shell on the surface, and tail friction μ_t was approximated by measurements of the C-legs of the robot on

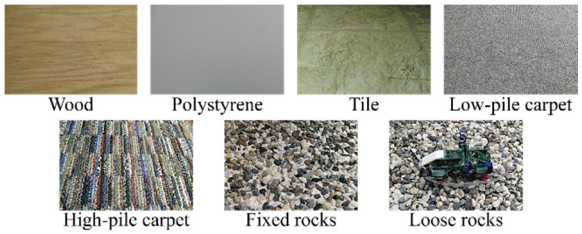


Fig. 16. Surfaces used in open-loop tail-assisted self-righting experiments.

TABLE II. TAILED RIGHTING PERFORMANCE ON VARIOUS SURFACES

Surface	**C.o.F.** μ_b/μ_t	**Model** t_r (ms)	**Experiment** t_r (ms)	**Successes**
Wood	0.4/0.5	235	274 ± 7.71	10/10
Polystyrene	0.3/0.3	217	264 ± 15.8	10/10
Tile	0.35/0.45	226	336 ± 51.7	10/10
Low-pile carpet	0.25/1	214*	259 ± 41.3	10/10
High-pile carpet	0.35/1	229*	256 ± 35.0	10/10
Loose rocks	0.35/1.2	231*	292 ± 65.8	9/10
Fixed rocks	0.35/1.2	231*	358 ± 138	8/10

the surface. The righting times from the dynamic simulation are also shown. Note that the model results with an asterisk had non-zero initial horizontal velocity—due to the high tail friction, the simulation executions from zero initial velocity resulted in a sticking tail with non-physical contact forces.

For the more uniform wood, polystyrene, and carpet surfaces, lower body friction and higher tail friction are correlated with faster righting times. The fastest average righting time was on high-pile carpet ($t_r = 256$ ms). For the wood, polystyrene, and low-pile carpet surfaces, the simulation matches the trend in righting times, but consistently underpredicts righting times by 39–47 ms.

The surfaces with more variation in roughness and appearance (tile, fixed rocks, loose rocks) had the highest average righting times with relatively large spreads across ten trials. The tile surface has millimeter-scale surface variations and the rock surfaces have centimeter-scale rocks either fixed in a single layer or loosely packed in multiple layers. Due to the variability of these surfaces, the tail can intermittently lose contact, which results in slower average righting times than predicted by the model. For the fixed and loose rock surfaces, the loss of tail effectiveness was sometimes so severe that the robot failed to right after a single tail swing.

E. Autonomous Self-righting Reflex on Unstructured Terrain

In addition to open-loop tests of tail-assisted self-righting, an autonomous self-righting reflex was demonstrated as VelociRoACH walked down a paper-covered obstacle with multiple step drops on the order of the robot's body length. A camera frame sequence of this experiment is shown in Fig. 17 and telemetry data is shown in Fig. 18. With the robot's legs following a bounding gait with a stride frequency of 4 Hz, a pre-programmed tail righting trajectory is activated during inversion. The trajectory periodically swings the tail over the top of the robot and is parameterized by a swing period of 0.4 s and a swing amplitude of 120° relative to the vertical tail position $\theta_t - \theta_b = 90°$. Using an onboard estimate of roll-pitch-yaw Euler angles (initialized with the accelerometer, then propagated forward with integrated gyroscope measurements), the tail PID controller is activated when the robot's body-fixed vertical vector $\mathbf{e_z}$ is aligned with the normalized gravity vector $\hat{\mathbf{g}}$. Using feedback on the tail position, the tail PID controller actuates the tail until the robot is upright.

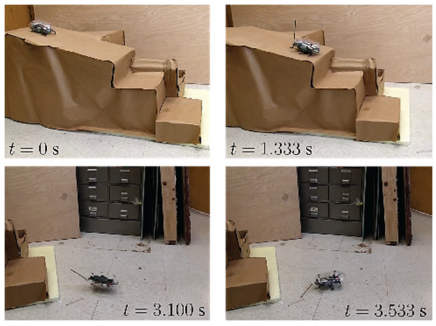


Fig. 17. Camera frame sequence of an autonomous tailed self-righting reflex after VelociRoACH walks off an obstacle with steep drop-offs.

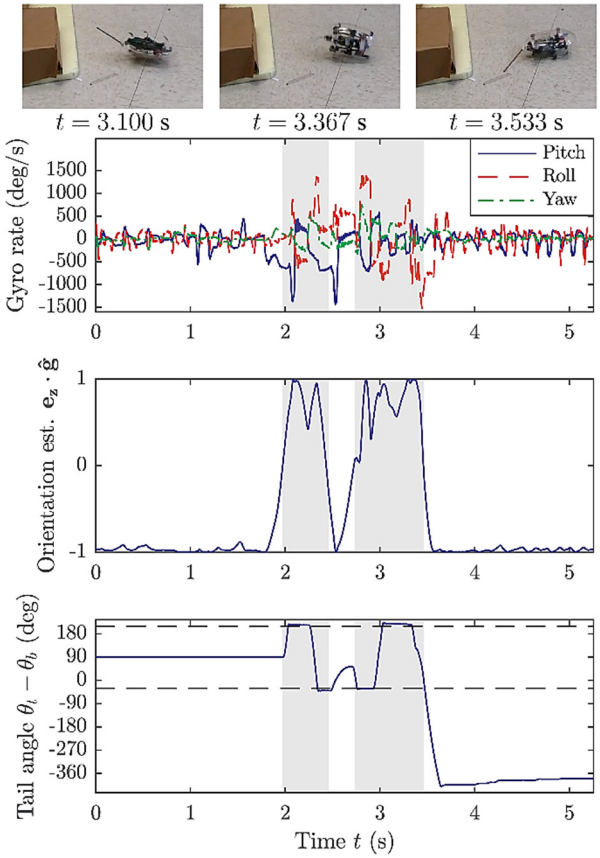


Fig. 18. Robot telemetry for an autonomous tailed self-righting reflex while walking down a step obstacle with a stride frequency of 4 Hz. The shaded time intervals are when the robot detects that it is inverted. In the tail angle plot, the dashed lines at −30° and 210° indicate the target tail angles.

As shown in Figs. 17 and 18, the autonomous self-righting reflex engages during locomotion whenever the robot's tumbling motion causes it to invert. In the second shaded time interval of inversion, the robot successfully self-rights at the bottom of the obstacle and walks forward.

V. Conclusions

The main conclusion from the tactile shell experiments is that a novel low-cost force/torque sensor can be used to measure the drag energy through narrow channels, which is correlated with an increase in the specific resistance of locomotion. However, some terrain contact can decrease the specific resistance of locomotion compared to free running. In future work, closed-loop strategies for steering to paths that minimize specific resistance will be developed, which is a potential transition of the results of this task to mobile robot technology.

The tail-assisted righting analysis gives design requirements and righting time predictions that can be extended to other MAST platforms that may benefit from a lightweight low-DOF tail for terrestrial self-righting. Experiments using a tailed VelociRoACH robot show that terrestrial self-righting can be consistently achieved with righting times as fast as 256 ms (with at least an 80% success rate after a single tail swing) on surfaces with different roughness scales and friction properties, including loose rocks. The demonstration of the autonomous tail-assisted self-righting reflex shows that VelociRoACH with tail is capable of detecting and recovering from inversion as it moves on complex terrain. The developed legged robot platforms with tactile shell and active tail push past the limits of small-scale locomotion by providing enhanced sensory feedback and greater mobility in heterogeneous terrain.

References

[1] C. Li, A. O. Pullin, D. W. Haldane, H. K. Lam, R. S. Fearing, and R. J. Full. "Terradynamically streamlined shapes in animals and robots enhance traversability through densely cluttered terrain." *Bioinspiration and Biomimetics*, vol. 10, no. 4, Jun. 2015.

[2] N. J. Cowan, E. J. Ma., M. Cutkosky, and R. J. Full. "A biologically inspired passive antenna for steering control of a running robot," in *Robotics Research: The Eleventh International Symposium*, 2005.

[3] J. Lee, S. N. Sponberg, O. Y. Loh, A. G. Lamperski, R. J. Full, and N. J. Cowan. "Templates and anchors for antenna-based wall following in cockroaches and robots." *IEEE Transactions on Robotics*, vol. 24, no. 1, Feb. 2008.

[4] J. T. Karras, D. W. Haldane, and R. S. Fearing. "Rapid-manufacturable hair sensor array for legged millirobots," in *IEEE/RSJ IROS*, Oct. 2012.

[5] T. Tsuji, Y. Kaneko, and S. Abe. "Whole-body force sensation by robot with outer shell," in *10th IEEE International Workshop on Advanced Motion Control*, 2008.

[6] T. Tsuji, Y. Kaneko, and S. Abe. "Whole-body force sensation by force sensor with shell-shaped end-effector." *IEEE Transactions on Industrial Electronics*, vol. 56, no. 5, May 2009.

[7] R. S. Fearing and T. O. Binford. "Using a cylindrical tactile sensor for determining curvature," in *IEEE ICRA*, Apr. 1988.

[8] A. Bicchi, J. K. Salisbury, and D. L. Brock. "Contact sensing from force measurements." *The International Journal of Robotics Research*, vol. 12, no. 3, Jun. 1993.

[9] O. Khatib. "A unified approach for motion and force control of robot manipulators: The operational space formulation." *IEEE Journal on Robotics and Automation*, vol. 3, no. 1, Feb. 1987.

[10] J. D. Goldberg and R. S. Fearing. "Force sensing shell using planar sensor for a miniature legged robot," in *IEEE/RSJ IROS*, Sep. 2015.

[11] T. Libby, T. Y. Moore, E. Chang-Siu, D. Li, D. J. Cohen, A. Jusufi, and R. J. Full. "Tail-assisted pitch control in lizards, robots and dinosaurs." *Nature*, vol. 481, no. 7830, Jan. 2012.

[12] E. Chang-Siu, T. Libby, M. Brown, R. J. Full, and M. Tomizuka. "A nonlinear feedback controller for aerial self-righting by a tailed robot," in *IEEE ICRA*, May 2013.

[13] T. Libby, A. M. Johnson, E. Chang-Siu, R. J. Full, and D. E. Koditschek. "Comparative design, scaling, and control of appendages for inertial reorientation." *IEEE Transactions on Robotics*, vol. 32, no. 6, Dec. 2016.

[14] G. Krummel, K. N. Kaipa, and S. K. Gupta. "A horseshoe crab inspired surf zone robot with righting capabilities," in *ASME IDETC*, Aug. 2014.

[15] C. Li, C. C. Kessens, A. Young, R. S. Fearing, and R. J. Full. "Cockroach-inspired winged robot reveals principles of ground-based dynamic self-righting," in *IEEE/RSJ IROS*, Oct. 2016.

[16] C. C. Kessens, D. C. Smith, and P. R. Osteen. "A framework for autonomous self-righting of a generic robot on sloped planar surfaces," in *IEEE ICRA*, May 2012.

[17] C. C. Kessens, C. T. Lennon, and J. Collins. "A metric for self-rightability and understanding its relationship to simple morphologies," in *IEEE/RSJ IROS*, Sep. 2014.

[18] D. W. Haldane, K. C. Peterson, F. L. G. Bermudez, and R. S. Fearing. "Animal-inspired design and aerodynamic stabilization of a hexapedal millirobot," in *IEEE ICRA*, May 2013.

[19] C. S. Casarez and R. S. Fearing. "Dynamic terrestrial self-righting with a minimal tail," in *IEEE/RSJ IROS*, Sep. 2017.

Dynamic Locomotion in Urban Environment

MCE-17-2.5

Justin K. Yim, Duncan W. Haldane, Mark M. Plecnik, Ronald S. Fearing
University of California, Berkeley
Berkeley, United States

***Abstract*—We have developed a one hundred gram monopedal jumping robot named Salto. Salto demonstrates the highest vertical jumping agility of an untethered, electrically actuated robot. It uses series elastic power modulation to achieve repeated, high-power, controlled jumps. With this novel method of jumping actuation, we demonstrate dynamic locomotion over varied terrain including the use of vertical surfaces to gain height through wall jumping.** *(Abstract)*

***Index Terms*—vertical jumping agility, series elastic power modulation, wall jump.** *(key words)*

I. Introduction

A. Background

Numerous jumping robots wind a spring with a small motor, and use an escapement mechanism to generate a high peak power instantaneously by releasing the stored energy. This strategy, termed the "catapult mechanism", is used by simple jumping insects such as fleas or grasshoppers [1] to great effect, and has been used by many robots. A 1.1 gram shape memory alloy robot has been made with the smart composite microstructures process using a catapult mechanism [2]. A 7 gram jumping robot [3] winds and releases a torsional spring to jump, and has a manually adjustable take off angle. The "Grillo" is a 15 gram jumping robot powered by a wind-up spring [4]. Larger jumping platforms are capable of directed jumps by changing the orientation of the jumping appendage prior to leaping. A cage was added to the 7 gram hopping robot to enable reorientation of the jumping appendage [5]. The "Jollbot" stores and releases strain energy which is stored in its protective shell [7]. The Sandia hopper [8] is driven by an orientable combustion powered piston. Several wheel driven robots with jumping appendages have been developed ([9], [10]). The robot developed by Tsukagoshi et al. was special because it could modulate the range of its jumps by varying the pressure in the pneumatic cylinder it used for jumping.

For the most part, jumping with an explosive release of energy precluded use of robotic limbs for anything but high-power jumping. Some control of the jumping direction has been accomplished by adjusting the orientation of the jumping appendage but fine-grained control of the jump is still lacking.

Our jumping robot also needs the capability to control its attitude in the flight phase; this subject has been well addressed by previous research, such as the falling cat robot of Yamafuji et al. 1992 which reorients in mid-drop. Attitude controllers for inertial bodies in free-fall have been well studied ([12], [13], [14], [15], [16], [17]). In addition, the use of inertial righting

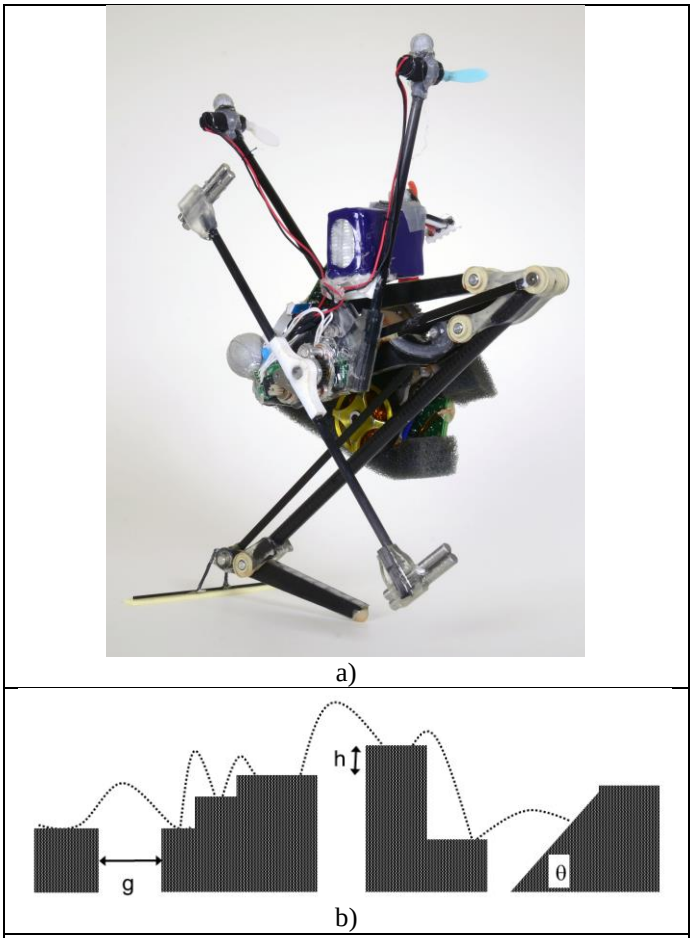


Fig. 1. a) Salto robot prototype b) Example terrain, including gaps, steps, and slopes.

by terrestrial animals has been studied in cats (Kane & Scher 1969), and more recently in Agama lizards in the lab of MAST collaborator Robert Full ([19], [20]). Based on the lizard findings, several robots were developed with one [21] and two degrees of freedom [22]. Recently another catapult mechanism jumper has been built [23], which has integrated a one degree of freedom tail to control pitch orientation [24].

Many robots have been built which can accomplish a small subset of the goals for the Capstone demonstration. These prior works serve as proof-of-concept platforms for the sub-systems integrated in Salto. The issue with previous jumping platforms is that the use of an explosive release mechanism precludes any notion of using that same jumping limb for other tasks such as highly controlled jumps and landings, landing on vertical

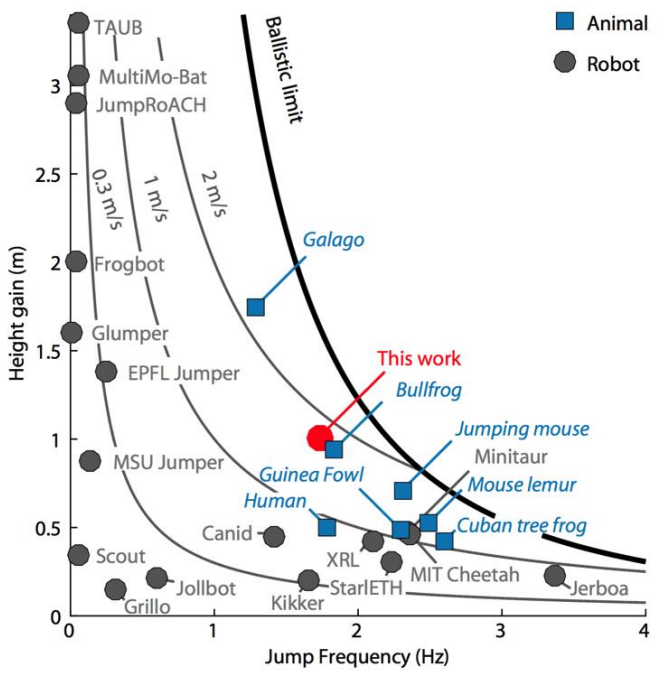


Fig. 2. Height gain and jump frequency for animal and untethered robotic systems.

surfaces, or locomotion in constrained environments. In the next section, we show that excellent jumping performance can be attained without resorting to a hard-to-control catapult mechanism.

To address the shortcomings of robots as agile locomotors as compared to animals, we first seek to establish an operational space and a metric upon which the agility of animal and robotic systems can be compared.

Recent work has proposed the capacity of a robot to increase its extrinsic energy in one jump, normalized by the platform mass, as a measure of its agility [25]. We use jump height as our first metric, shown in Fig. 2, which is simply this measure further normalized by the gravitational constant. To characterize the temporal characteristics of these agile motions, the number of times a robot can perform a maneuver in succession (termed endurance) was considered [25]. A robot that cannot immediately perform a jump after landing is fixed to an endurance of n = 1, and n → ∞ corresponds to an unlimited amount of jumps without overheating its actuators.

We seek to slightly extend these previously proposed measures, to allow comparison of a larger range of animal and robotic jumping systems, without coupling the metric to a specific actuation strategy. Therefore we consider the rate at which systems can jump as the second dimension of the operational space of agile systems. This metric is presented as jump frequency, on the horizontal axis of Fig. 1. The jump frequency is calculated as:

$$f = 1/(t_s + t_f)$$

where f is the jump frequency, t_f is the flight time of the jumper from when it lifts off from the ground to the apogee of its leap, and t_s is the stance time of the system during which it applies power to increase its extrinsic energy. This value is calculated for a single jumping maneuver, and corresponds to the fastest possible rate a robot could chain together a series of vertical leaps, as would be used for running up stairs.

By defining the operation space in this way, we can compare the agility of jumping systems with a diverse range of mechanisms including wind-up jumping robots, muscle-powered accelerations in animals, and robots capable of continuous locomotion. We seek to inform the design of controllable, agile robotic systems capable of field operation, so not considered in the comparison are robots which can draw an unlimited amount of power via a tether, or robots that uncontrollably generate jumping energy via a combustion chamber.

For a joint metric over this space, we define agility as the product of jump frequency and jump height. Agility has units of m/s, and corresponds to the maximum achievable average vertical velocity of the jumping system while performing repeated jumps. Gravity imposes a kinetic limit on the total agility of a system. Salto achieves a vertical jumping agility of 1.75 m/s, the highest of any untethered, electrically actuated robot.

Material in the introduction, technical approach, and results sections of this report has been previously published in [26], [27], [28], and [29].

B. Research Objectives

We aim to develop a robotic platform weighing approximately one hundred grams, which is capable of controlled long-range jumps, dynamic transitions on-to and off-of walls, as well as locomotion over rough terrain. These diverse capabilities place multifunctional demands on robotic appendages, which have in the past precluded their comprehensive integration on a single robotic platform. Long-range jumps require high peak power output; controlling the range and direction of these jumps and controlled transitions require fine-grained control over limb trajectories and force production. By understanding how previous robots have accomplished or failed these tasks, and how animals accomplish them all simultaneously, we show how they can be integrated on a single robotic platform. The capstone project involves research in mechanics, control, and design combined with system engineering of a complete integrated robot system, capable of untethered locomotion, with assisted autonomy (such as using Vicon for state estimation and externally provided terrain information).

C. Technical Approach Summary

Based on [30], we take inspiration from the Galago senegalensis in design of Salto's series elastic power modulating leg. Using series elastic actuation in conjunction with a variable mechanical advantage leg that modulates force output, series elastic power modulation can produce higher peak power than the actuator alone could deliver. In order to fulfill the multifunctional demands for this leg, we use a multi-step design process. We begin with kinematic synthesis for design exploration of potential six-bar linkages, followed by kinematic tuning by optimization. Resulting designs are vetted

in physical simulation. Physical tests of designs inform physical simulation parameters. The resulting robot weighs approximately one hundred grams, is capable of vertical jumps up to a meter from the ground, and can use vertical surfaces to increase its height through wall jumping. We demonstrate an example of a new class of dynamic maneuvers made possible by high vertical jumping agility by performing a wall jump. Furthermore, we demonstrate the robot's ability repeatedly execute high jumps with the addition of aerodynamic thrusters for stabilization.

II. Related Work

A. Synergy with Existing Works and External Collaborations

Salto can greatly benefit from progress made in many research areas in the MAST project. Existing works from past/completed MAST tasks include:

MAST Task	Dependency/Interactions
MCE-2.5 (FY08-15) Foot and ankle design for dynamic locomotion	Will use results on appendages with claws, engagement, release, and tuned dynamic behavior
MCE 2.4 (FY08-14) Bio-inspired Crawling, Running, Climbing Robots	Follow-the-leader using Heimann IR imager (as shown in 2013)

Continuing internal and external collaborations within current MAST tasks include:

MAST Task	Dependency/Interactions	PI
MCE 2.1 Ground interaction	Will use existing and new results on loose media terramechanics for possible running/jumping on loose media	Goldman, Georgia Tech
MCE 2.2 Bio-inspired principles of extended terrain mobility	Will use existing and new insights on bioinspired appendage design	Full, Berkeley
MCE 2.3 Bio-inspired Principles of Transitions	New insights into transition modes, for moving between horizontal and vertical surfaces	Full, Berkeley
MCE 2.4 (FY16, 17) Exoskeleton for Millirobot Locomotion in Rugged Terrain	Salto will need robust exoskeleton for loose media, falls, body righting	Fearing, Berkeley

B. Technology Transitions Summary

For the capstone demonstration, we will demonstrate Salto, a palm-sized legged robot that is able to locomote in challenging terrain, including jumping over gaps, climbing steps, and climbing steep slopes. Using an inertial tail, the Salto can demonstrate high maneuverability in the vertical plane (jumping over or under obstacles), and control landing attitude. We will assume that locations of terrain features are provided ahead of time to remove dependence on mapping from other capstone tasks (although collaboration with other MAST capstones which combine ground and aerial mobility would be welcome). (Landing on a target would be a great use for optical flow sensing, which is a potential collaboration with ARL.) For continuation of this work, a proposal is being submitted to Dr. Sam Stanton of ARO, entitled: ``3D Saltatorial Locomotion on Uneven Terrain''.

C. Technical Approach

The multifunctional demands of controllable jumps and dynamic maneuvers preclude the use of a catapult mechanism to power the robot. To enable large peak power output without relying on a single explosive release of energy, we use a power-dense brushless motor in conjunction with a series elastic element. Small hobby-grade outrunner motors can provide 2.9 kW/kg in a twelve-gram package (e.g. www.hobbyking.com). The outrunner form factor maximizes the gap radius of the motor, minimizing the required gear ratio and thus the reflected inertia of the motor's rotor. This reduces the time constant of the actuator, allowing power to be applied more quickly. An elastic element placed in series with the actuator can amplify its power output, increase the time it can operate near its peak power point, and enable force control ([31], [32]). Jumping animals such as the ~250 gram *G. senegalensis* (also known as lesser bush-baby) use tendon structures and muscular aponeurosis as series-elastic elements to greatly amplify their power output while jumping [30]. The maximum power density of the fast glycolytic muscle fibers that power the bush-baby's jump is 0.38 kW/kg [33], over seven times less than the brushless motor. However the animal is still capable of leaping 1.74 meters vertically due to its series elastic structures [30]. We use a high-power brushless motor to power the jump. With a stance time of 120 milliseconds, Salto achieves a maximum energy input of 2J. With the choice of series elastic actuators driven by high-power brushless motors, we have shown that a robot can produce agile jumps without resorting to specialized mechanisms that can only release energy explosively. This actuator choice enables a highly diverse range of jumps, and controllers to regulate force application during stance.

The purpose of the leg mechanism is to transform spring torque into a specified ground reaction force that acts through the center of mass of the prototype. The space of engineering solutions that we explored was mechanical linkages composed of revolute joints. The challenge of this design paradigm is that the parameters are embedded in a highly nonlinear geometric space. We approach this challenge with a two-stage method starting with an initial design exploration, followed by a kinematic tuning ([28], [34]).

Design exploration determines locations in the design space where required behaviors are coarsely achieved. This was accomplished by formulating polynomial synthesis equations

and solving them with the homotopy continuation solver Bertini. Kinematic tuning begins with one of these locations and then optimizes to satisfy all requirements to the granularity desired. To begin design exploration, we identified a set of primary required behaviors: (i) the mechanism constrains the foot point to a straight line, the line of action; (ii) foot stroke is long; (iii) pivots are located above the foot with the input near the line of action; and (iv) link lengths are compact.

Using the design procedure outlined by Plecnik and McCarthy [34], an atlas of Stephenson six-bar linkages was generated, marking locations of the design space to explore. The design that best produced the primary behaviors was advanced to the kinematic tuning phase to include an extended list of behaviors: (v) the input link attached to the series-elastic actuator rotates over a large range; (vi) a low starting MA; (vii) MA defines a constant force at the foot during the rest of stroke; and (viii) moments exerted on the body are minimized. The optimization was performed using the InteriorPoint method within Mathematica's FindMinimum function.

We dynamically simulated (using approximations of link mass and inertia) a maximal jump to determine the success of the synthesis approach and the magnitude of internal reaction forces and then designed a spring-driven prototype. Experiments with this prototype validated the design of the energy-storing spring, but inaccurate estimations of link masses caused the jumping motion to generate substantial angular momentum. Accurate mass models from the prototype showed that our six-bar design was difficult to balance in order to satisfy (viii) above. Two more links were added to the six-bar linkage, and optimization was performed again to arrive at the final design shown in Fig. 3B. A spring-driven prototype of this refined linkage verified that the design was balanced.

Simulations show that the leg linkage produces internal forces up to 210 N. To minimize mass, as much of the linkage as possible is made from carbon fiber composite materials. Links in compression or bending are milled from prefabricated honeycomb core carbon fiber composite plates (XS-LP-39, ACP Composites) using a Dremel 569D bit installed in an enclosed precision desktop mill (OtherMill, Other Machine Co.). The torsional stiffness of these sandwich composite members is increased by adhering a 0.005-inch-thick FR-4 fiberglass sheet cut on a 45° bias (1331T23, McMaster-Carr),

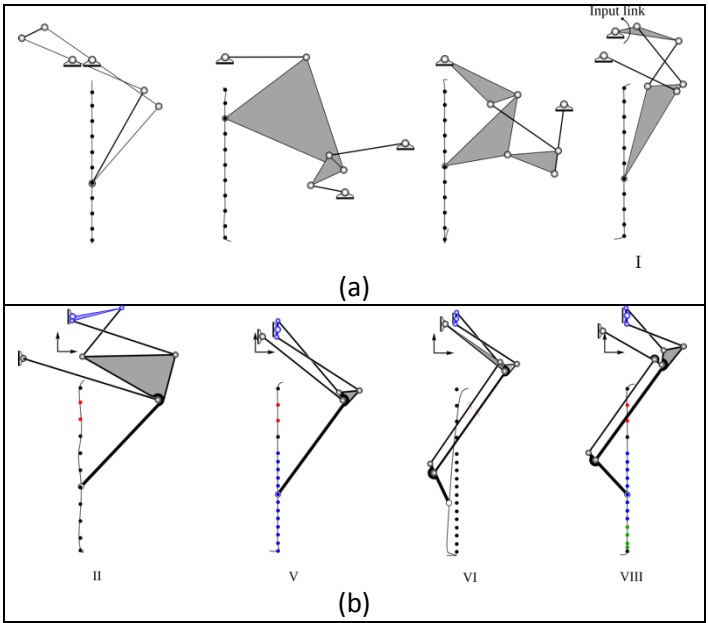


Fig. 3. a) Selected candidate six-bars from the design exploration atlas. b) Design iterations in kinematic tuning

across the width of the beams with cyanoacrylate adhesive (Loctite 495). The two links that carry purely tensile loads are fabricated as tie rods with unidirectional pultruded carbon fiber flat stock (2153T21, McMaster-Carr). The revolute joints of the mechanism are formed by a precision ground 2024 aluminum rod (9062K24, McMaster-Carr) riding in polymer bushings (JFM-0304-05, igus Inc.). The bushings and shafts are integrated with the structural carbon fiber using molded polyurethane components (IE- 3075, Innovative Polymers). A two-stage casting process produces these parts; wherein a solid wax mold is cut using OtherMill, a silicone mold is taken from this positive, and the final polyurethane parts are cast from the silicone mold. This fabrication paradigm resulted in a linkage weighing only 11 g that can withstand the internally generated forces and that did not fail during experiments.

The torsional spring is cut from continuous-dip processed latex (5234K681, McMaster-Carr) into a conic section using a 3D printed jig. Spring material properties and a model for leg friction are identified by fixing the motor position and cyclically compressing the leg in an Instron material testing machine.

The brushless motor starts as a Scorpion S-1804-1650KV winding kit obtained from the factory, wound with 48 turns of 30 AWG (American wire gauge) magnet wire. We statically verified the torque constant. The custom wound stator is mounted in a custom 25:1 gearbox for a total mass of 25.29 g. An AS5047P magnetic encoder calibrated to rotor position generates commutation signals for a custom brushless DC motor driver developed in-house. The reference is generated by an ImageProc robot control board [35] (printed circuit board: https://github.com/biomimetics/imageproc_pcb, firmware: https://github.com/dhaldane/roach). A second AS5048B magnetic encoder measures the angle of the lower ternary link; the angular deflection of the spring is estimated with an inverse-kinematics look-up table. The system is powered by a three-cell 11.1-V 180 mA·hour lithium polymer battery (Turnigy nano-tech).

The control board records telemetry (six-axis inertial measurement unit, estimated body angle, leg position, and motor position) at 1000 Hz. It is controlled from a laptop via an onboard ZigBee radio interface. Contact with terrain is detected by spring displacement exceeding a threshold or the rotational rate of the robot exceeding a threshold. This detector ensemble is more robust against variations in the angle of incidence than either detector in isolation.

For stable and repeatable locomotion, it is important to control the orientation of the robot as it makes transitions. The leg angle at touchdown is important for stability [6], and the orientation of the robot is critical as it approaches a wall to jump [36]. We use an inertial tail to control the attitude of the robotic platform while it is in the air. Another option would be to use momentum wheels, but tails serve better as multi-functional appendages. They can be used for rapid terrestrial righting behaviors (Unpublished observations from Agama agama), assist in robotic climbing ([37], [38]), and be used for balance during stance [39].

The tail is simple in design. Previous work [40] has shown that the tails most effective for inertial righting (per unit mass) are point masses on weightless beams.

Additional developments to Salto for repeated spatial jumping resulted in a modified robotic platform named Salto-1P. Due to Salto's long flight phase and extremely short stance phase (a characteristic of its high vertical jumping agility), we selected small aerodynamic thrusters for stabilizing the robot's roll and yaw rotations out of the sagittal plane. These thrusters are commercial mini- quadcopter (Cheerson CX-10) propeller blades mounted in a 'V' configuration, with moment arms about the center of mass that prioritize control over roll angle (80mm) over yaw angle (40mm). Roll torque is created by driving both thrusters in the same direction; yaw torque is generated with a differential motor command.

To investigate performance benefits brought by improved ground interaction, we designed a gripper mechanism for Salto that engages angled spines and is kinematically constrained to engage/disengage with leg crouch/extension. To coordinate leg configuration to gripper engagement we used the relative angle between the tibia and tarsus links. Angle ϕ is plotted in Fig. 4a along with the desired coordination between leg extension and gripper abduction angle $\Delta\Theta$. The gripper moves to an engaged configuration as the robot begins to crouch, then remains engaged for all other leg configurations, see Fig. 4a. Moreover, gripper orientation must be decoupled from engagement motion, allowing the contact angle with a surface to define this additional degree-of-freedom.

The gripper is made up of three separate components: the capstan, the string system, and the feet. The capstan is 3D printed and attached to the base of the tibia. The string system consists of the pins, string tightener, and an aramid fiber string. The small pins are to route the string through the gripper system, the linear actuator pin pulls the feet into an "engaged" configuration, and there is a central pin that the feet rotate around. There are three separate feet for force balancing. The top half of a foot consists of a milled out polypropylene piece with a hole to connect to the linear actuator pin and a living hinge that acts as a joint. The polypropylene is glued to a 3D printed piece with holes to direct the exact angle of the spines, which are made from steel pins.

We performed tests of the Coefficient of Engagement (CoE), which is the ratio of a maximal statically generated normal load or tangential load to disengage from a surface to the impact force that enabled that load. We tested the gripper's performance against a rubber ball foot, fixed spines, and a gripper covered in rubber. We also tested it on cork, sand paper, rigid foam, and acrylic. We later used these CoE's in our dynamic model to predict at which heights the robot would be able to jump depending on the surface and foot type. Next we jumped a robot with the various feet on different surfaces at specific angles. Video stills from two example jumps are shown in Fig. 4b where the spiny gripper jumps at 60° and the rubber ball slips at 30°. We compared these jump tests to the dynamic model and found that it closely matched (Fig. 4c).

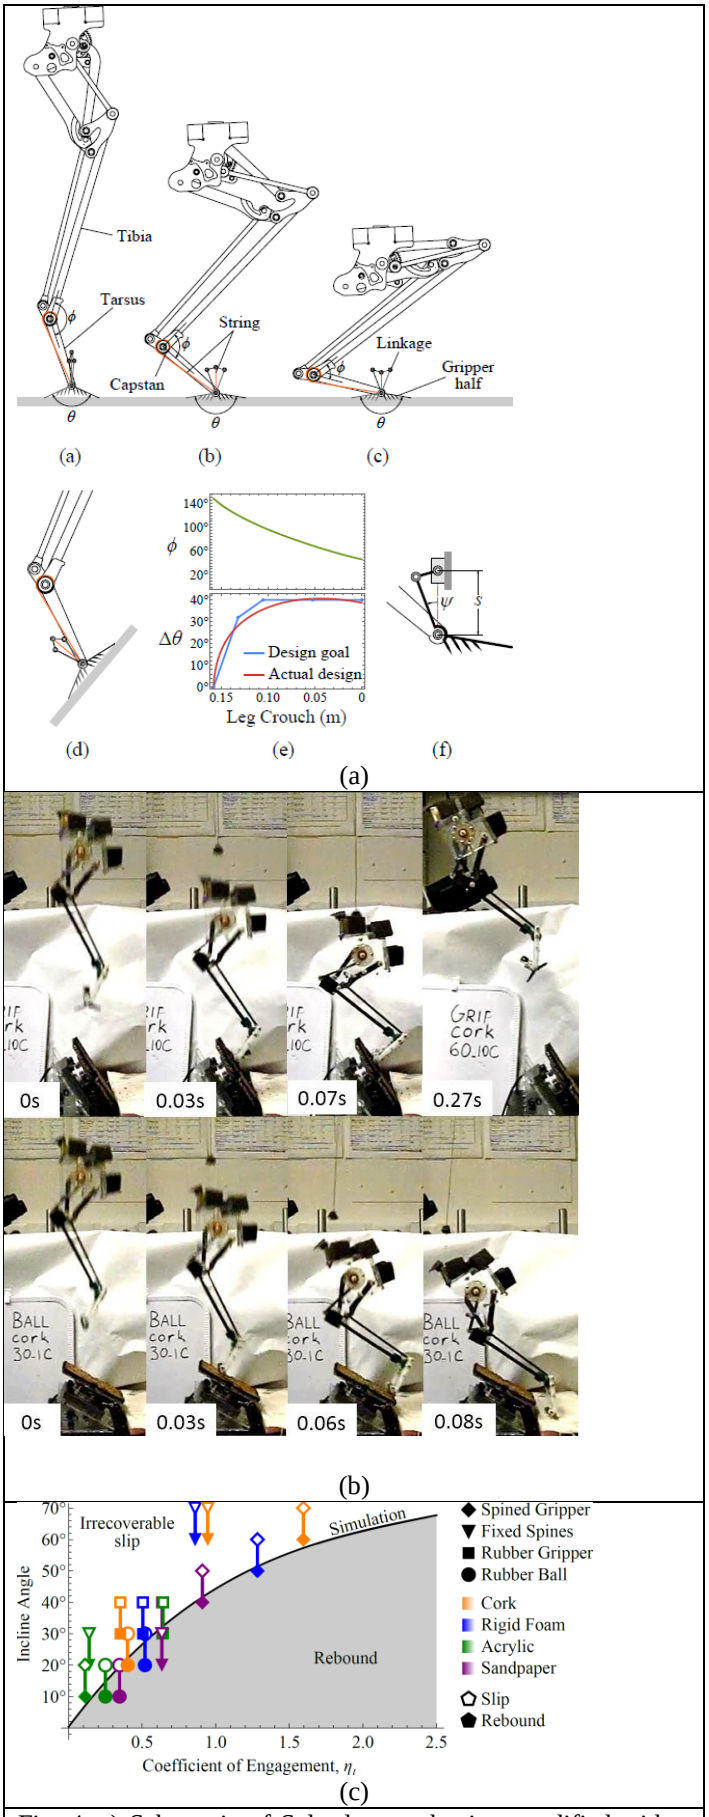


Fig. 4. a) Schematic of Salto leg mechanism modified with a kinematically coordinated gripper. b) Comparison of jump tests and the dynamic model. c) Jump testing of Robot with spined gripper (top) and rubber ball (bottom).

III. Results

A. Research Accomplishments

The Salto leg mechanism simultaneously accomplishes multiple functions. In order to compare our work to prior literature, we list the following:

A. Actuator torque/force is transformed into a vertical ground reaction force (GRF) that exceeds the robot's weight.

B. The mechanism reduces the motor and gearing requirement.

C. Desired take-off velocity is achieved with minimal peak accelerations.

D. The robot jumps with near-zero angular velocity.

E. Transient power output of a series elastic actuator is modulated to increase jump energy.

By simultaneously producing all of the above functions, our design is set apart from the myriad of novel jumping robots produced over the past few decades.

The result of the design methodology is a monopedal jumper composed of a linkage that enables power modulation, a series-elastic actuator, and an inertial tail for attitude control. Salto was designed to minimize mass so that less energy is required per unit height gain and to improve robustness against high-energy collisions created by agile locomotion. It jumps in free space, with neither tether nor planarizing boom, both of which might interfere with agile motions.

We use a planar eight-bar revolute linkage to encode a Mechanical Advantage (MA) profile enabling power modulation, as well as fundamental attributes that decouple leaping motions from rotational dynamics. We conceptualize the MA (shown in Fig. 5) as piecewise constant for the purpose of designing an SE+MA jumping appendage. The profile starts low at 2.49 N/Nm in the energy storage phase and increases to an average of 26.90 N/Nm in the energy return phase, with a G_R value of 10.8. To limit damaging peak forces during energy release, we specify that the MA function produce a constant force when the leg mechanism leaves the energy storage phase, reducing the peak acceleration (~20g) by a factor of 2, as compared with a constant-MA case.

We also specify that extension of the leg generate no body rotation to enforce decoupled dynamics. This was accomplished by specifying that the foot move along a straight line that passes through the center of gravity of the linkage and the robot body.

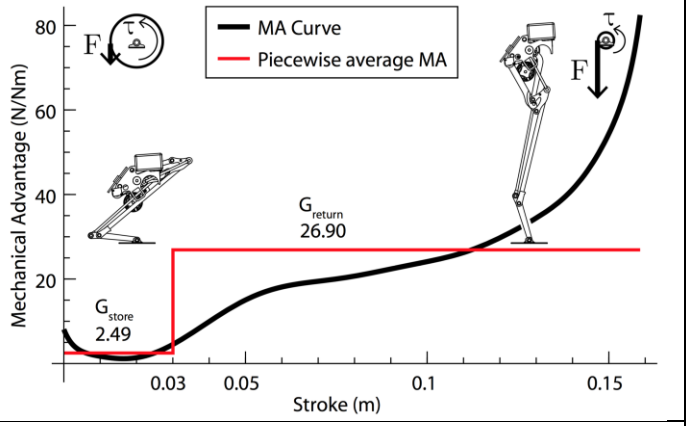


Fig. 5. Mechanical advantage as a function of leg extension. Average values of G_{store} and G_{return} for energy storage and return phases are shown. Cartoons show the leg mechanism fully crouched and fully extended.

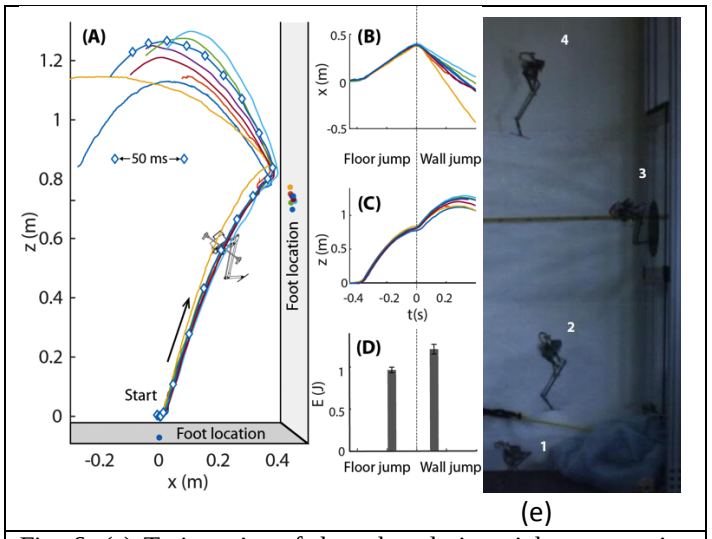


Fig. 6. (a) Trajectories of the robot during eight consecutive wall-jump experiments (b) horizontal component of trajectories vs time (c) vertical component of trajectories vs time (d) Energetic state of center of mass before and after wall-jump. (e) Composite of stills from a wall jump experiment.

A series-elastic actuator composed of a torsional spring and geared brushless DC motor drives the linkage. For energy density and compactness, the conic-sectioned torsional latex spring developed for the SEA snake (Series Elastic Actuator snake) [41] was used. This design fully loads the elastic material and leverages the latex's higher energy storage per unit mass compared with other common spring materials such as composites or steel.

To anchor the meaning of vertical jumping agility in the context of locomotion, we demonstrate a wall jump with the Salto robotic prototype shown in Fig. 6. In this experiment, the robot starts in a stable configuration on the floor, uses the tail to orient itself toward the wall, jumps, reorients, and then jumps again off the wall. Parameters for this jump were tuned to maximize the height gained off the wall. The motor and spring power are shown in Fig. 7a for the starting jump and in Fig. 7b for the wall contact phase.

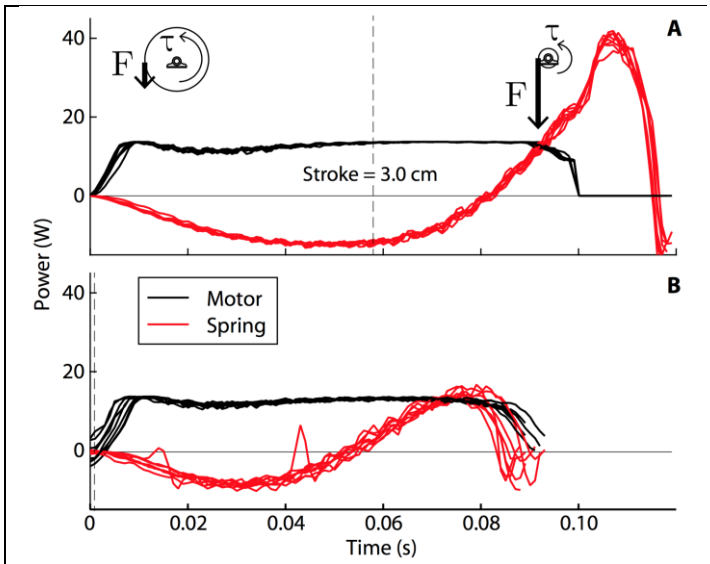


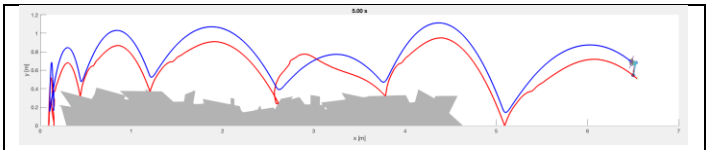
Fig. 7. Spring and motor power for (a)the starting jump of the wall jump maneuver and (b) the wall contact phase (N = 8). The dashed line indicates where the leg moves from the storage to return MA regions (shown in Fig. 4)

Figure 7a shows motor and spring power as a function of time for the robot jumping from its fully crouched position. This figure shows the effect of the SE+MA appendage on power modulation. Eight consecutive trials are shown on this plot, with the time axis starting with the onset of motor activation. The leg stays in the low-MA region (shown in Fig. 5) until 0.06 s into the stroke, loading energy into the spring. After this point, the MA increases and the energy in the spring is returned to accelerate the robot. The average (N = 8) peak spring and motor power were 40.3 ± 1.54 and 13.7 ± 0.0 W, respectively, showing that the spring provided 2.94 times greater power than the motor alone can produce. These data also show the consistency of the power-modulating behavior when the platform starts from rest.

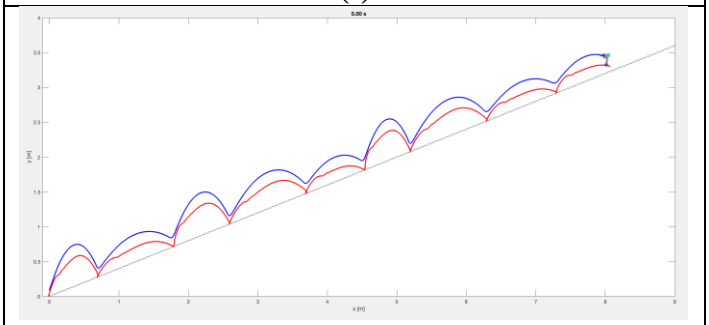
When jumping off the wall, the power-modulating effect is less prominent because contact is made with the leg extended out of the energy storage region, as shown in Fig. 7b. Integrating the motor power shows that an average of 1.22 ± 0.01 J was delivered during the starting jump, and 1.04 ± 0.03 J was delivered during the wall contact phase. Variation in initial conditions when the robot contacts the wall caused this phase to be more variable than the starting jump.

The robot attained an average height gain of 1.21 ± 0.065 m for the wall jump, reaching an average of 0.4 m higher than the foothold location on the wall. The maximum height increase was 1.29 m, 0.28 m higher than the maximum achievable were the platform jumping from the floor.

In addition to the short-duration wall jump maneuvers, continuous locomotion on flat ground was investigated in a Vicon motion capture environment. Fig. 8 shows motion



(a)

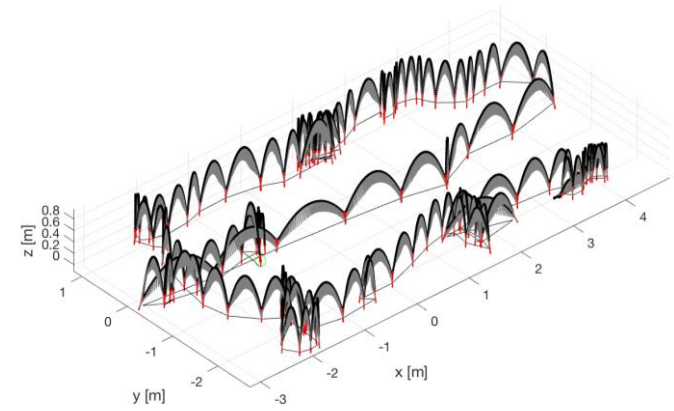


(b)

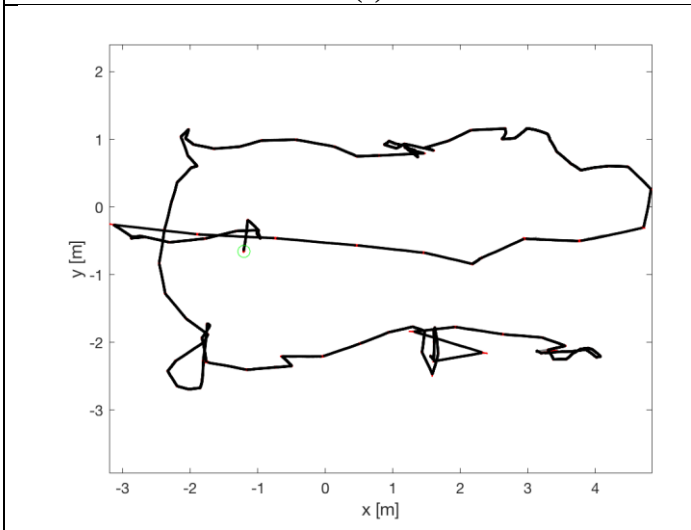


(c)

Fig. 9. a) Physical 3D spatial jumping experimental trajectory. b) Trajectory overhead view. c) Video still of Salto-1P from a similar experiment.



(a)



(b)



(c)

Fig. 8. a) Vertical planar simulation of Salto on rubble. b) Ramp climbing; the ramp is toothed to prevent foot sliding. c) Video still of Salto-1P (upper left) jumping over discontinuous uneven surfaces.

capture trajectories and a video still from these experiments. As in the wall jump, most motions occurred in the sagittal plane with pitch orientation controlled by the inertial tail. The aerodynamic thrusters stabilized roll and yaw and produced slower maneuvers out of plane like pirouettes around the vertical axis. The robot demonstrated up to 172 consecutive jumps and travelled over 25 m horizontally during a single run. It demonstrated forwards running at up to 3.6 m/s and a long jump over 2.0 m in length. Fig. 8a and 8b show a 25 m trajectory that includes running at various speeds forwards and backwards, turning in place, and coordinated turns.

Beyond flat ground, we investigated jumping over discontinuous surfaces of various slopes and stiffnesses. First, we simulated jumping in a custom Matlab dynamic simulator shown in Fig. 9a and 9b. Simple linear control laws similar to [44] produced stable jumping over simulated rubble and up slopes since Salto's high vertical jumping agility produces easy addition of energy. Note that the controller recovers from the foot sliding near the center of the rubble pile. Following simulated experiments, we conducted physical jumping experiments in the Vicon motion capture environment shown in Fig 9c. Improvements in sensing and control to improve the robustness of jumping on complex terrain are the subject of ongoing work.

B. Metrics for Performance Evaluation

IV. Sub-Task	Metric	Basic Goal	Stretch Goal
Controlled Jumping	height	1 m	2 m
"	range	0.5 m	4 m
"	frequency	1.5 Hz	5 Hz
"	vertical jumping agility	1.5 m/s	3 m/s
Re-orientation	time	400 ms	(200 ms)
"	min tail mass for reorientation	15 gm	5 gm
Floor to Wall Transition	safe landing velocity range, and safe landing angle	normal: 1.2-1.8 m/s 0-60 degrees	normal: 1-2 m/s 0-90 degrees
Wall to Floor Transition	velocities and angles for launching	0-2 m/s, 0-90 vertical, +- 10 orthogonal	0-4 m/s, 0-90 vertical, +- 10 orthogonal
Wall to Wall Jumping	Span of distances and orientations between walls for successful jumping	0.5-1 m, 0-60 degrees	0.5-2m, 0-90 degrees

Note that the vertical jumping agility performance metric is not equivalent to the product of the jump height and jump frequency since higher jump frequency could be achieved in jumps that are not to apex. This occurs in situations like stair climbing in which the robot's vertical velocity never reaches zero. Jumping to 2 m at the ballistic limit in which the robot has no stance time corresponds to a vertical jumping agility of 3.18.

V. Conclusions and Further Work

A. Potential Technology Transitions

We have filed a provisional patent for Salto's leg linkage with the University of California, Berkeley Office of Intellectual Property & Industry Research Alliances.

References

[1] R. M. Alexander, "Leg design and jumping technique for humans, other vertebrates and insects.," *Philos. Trans. R. Soc. Lond. B. Biol. Sci.*, vol. 347, no. 1321, pp. 235–48, Feb. 1995.

[2] M. Noh, S.-W. Kim, S. An, J.-S. Koh, and K. Cho, "Flea-Inspired Catapult Mechanism for Miniature Jumping Robots," *IEEE Trans. Robot.*, vol. 28, no. 5, pp. 1007–1018, Oct. 2012.

[3] M. Kovac, M. Fuchs, A. Guignard, J.-C. Zufferey, and D. Floreano, "A miniature 7g jumping robot," *IEEE/RSJ Int. Conf. Robot. Autom.*, no. c, pp. 373–378, May 2008.

[4] U. Scarfoglierol, C. Stefanini, and P. Dariol, "Design and Development of the Long-Jumping ' Grillo ' Mini Robot," *IEEE Int. Conf. Robot. Autom.*, no. April, pp. 10–14, 2007.

[5] M. Kovač, M. Schlegel, J.-C. Zufferey, and D. Floreano, "Steerable miniature jumping robot," *Auton. Robots*, vol. 28, no. 3, pp. 295–306, Dec. 2009.

[6] B. Andrews, B. Miller, J. Schmitt, and J. E. Clark, "Running over unknown rough terrain with a one-legged planar robot.," *Bioinspir. Biomim.*, vol. 6, no. 2, p. 026009, Jun. 2011.

[7] R. Armour, K. Paskins, A. Bowyer, J. Vincent, W. Megill, and R. Bomphrey, "Jumping robots: a biomimetic solution to locomotion across rough terrain.," *Bioinspir. Biomim.*, vol. 2, no. 3, pp. S65–82, Sep. 2007.

[8] P. Weiss, "Hop … Hop … Hopbots !," *Science News*, vol. 159, 2001.

[9] S. A. Stoeter, P. E. Rybski, M. Gini, and N. Papanikolopoulos, "Autonomous Stair-Hopping with Scout Robots," *IEEE/RJS Int. Conf. Intell. Robot. Sys.* no. October, pp. 721–726, 2002.

[10] H. Tsukagoshi, M. Sasaki, A. Kitagawa, and T. Tanaka, "Design of a Higher Jumping Rescue Robot with the Optimized Pneumatic Drive," Proc. *IEEE Conf. Robot. Autom.* no. April, pp. 76–83, 2005.

[11] K. Yamafuji, T. Kobayashi, and T. Kawamura. "Elucidation of twisting motion of a falling cat and its realization by a robot." *J. Robot. Soc. Japan* 10.5 (1992): 648-655.

[12] K. Yoshida and Y. Umetani, "Control of Space Free-Flying Robot," *IEEE Conf. Decis. Control*, no. December 1980, pp. 4–9, 1990.

[13] E. Papadopoulos and S. Dubowsky, "On the nature of control algorithms for free-floating space manipulators," *IEEE Trans. Robot.*, 1991.

[14] C. Fernandes, L. Gumits, and Z. Li, "Near-Optimal Nonholonomic Motion Planning for a System of Coupled Rigid Bodies," *IEEE Trans. Automat. Contr.*, vol. 39, no. 3, 1994.

[15] C. Rui, I. V Kolmanovsky, and N. H. McClamroch, "Nonlinear Attitude and Shape Control of Spacecraft with Articulated Appendages and Reaction Wheels," *IEEE Trans. Automat. Contr.*, vol. 45, no. 8, pp. 1455–1469, 2000.

[16] Y.-L. Yang, P. C. Chao, and C.-K. Sung, "Landing Posture Control via Trajectory Planning for a Generalized Twin-body System," *Int. Conf. Auton. Robot. Agents*, pp. 97–102, 2009.

[17] T. Kai, "A Model Predictive Control Approach to Attitude Stabilization and Trajectory Tracking Control of a 3D Universal Joint Space Robot with an Initial Angular Momentum," *IEEE Conf. Decis. Control*, Dec. 2011.

[18] T. R. Kane and M. P. Scher, "A Dynamical Explanation of the Falling Cat Phenomenon," *Int. Jnl. Solids Struct.*, vol. 5, 1969.

[19] A. Jusufi, D. I. Goldman, S. Revzen, and R. J. Full, "Active tails enhance arboreal acrobatics in geckos.," Proc. *Natl. Acad. Sci. U. S. A.*, vol. 105, no. 11, pp. 4215–9, Mar. 2008.

[20] T. Libby, T. Y. Moore, E. Chang-Siu, D. Li, D. J. Cohen, A. Jusufi, and R. J. Full, "Tail-assisted pitch control in lizards, robots and dinosaurs.," *Nature*, vol. 481, no. 7380, pp. 181–4, Jan. 2012.

[21] E. Chang-Siu, T. Libby, M. Tomizuka, and R. J. Full, "A lizard-inspired active tail enables rapid maneuvers and dynamic stabilization in a terrestrial robot," *IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, pp. 1887–1894, Sep. 2011.

[22] E. Chang-siu, T. Libby, M. Brown, M. Tomizuka, and I. Introduction, "A nonlinear feedback controller for aerial self-righting by a tailed robot," 2013 *IEEE Int. Conf. Robot. Autom.*, May 2013.

[23] J. Zhao, N. Xi, B. Gao, M. W. Mutka, and L. Xiao, "Development of a controllable and continuous jumping robot," 2011 *IEEE Int. Conf. Robot. Autom.*, pp. 4614–4619, May 2011.

[24] J. Zhao, T. Zhao, and N. Xi, "Controlling aerial maneuvering of a miniature jumping robot using its tail," *IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, pp. 3802–3807, 2013.

[25] J. M. Duperret, G. D. Kenneally, J. L. Pusey, and D. E. Koditschek, "Towards a comparative measure of legged agility," in *Springer Tracts Adv. Robot.*, 2016, vol. 109, pp. 3–16.

[26] D. W. Haldane, M. M. Plecnik, J. K. Yim, R. S. Fearing, "A power modulating leg for monopedal hopping," *IEEE/RJS Int. Conf. Robot. Autom.*, 2016.

[27] D. W. Haldane, M. M. Plecnik, J. K. Yim, R. S. Fearing, "Robotic vertical jumping agility via series-elastic power modulation," *Science Robotics* Dec. 2016, pp.

[28] M. M. Plecnik, D. W. Haldane, J. K. Yim, R. S. Fearing, "Design exploration and kinematic tuning of a power modulating jumping monopod." *J. Mech. Robot.*, 10.1115/1.4035117 (2016).

[29] D. W. Haldane, J. K. Yim, R. S. Fearing, "Repetitive extreme-acceleration (14-g) spatial jumping with Salto-1P" *IEEE/RJS Int. Conf. Robot. Autom.*, 2017

[30] P. Aerts, "Vertical jumping in Galago senegalensis: the quest for an obligate mechanical power amplifier," *Trans. R. Soc.*, no. May 1997, pp. 1607–1620, 1998.

[31] D. Paluska and H. Herr, "The effect of series elasticity on actuator power and work output: Implications for robotic and prosthetic joint design," *Rob. Auton. Syst.*, vol. 54, no. 8, pp. 667–673, Aug. 2006.

[32] G. A. Pratt and M. M. Williamson, "Series elastic actuators," *IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, vol. 1, pp. 399–406, 1995.

[33] B. M. Nigg and W. Herzog, *Biomechanics of the musculo-skeletal system*, 2nd ed. Wiley, 1999.

[34] M. M. Plecnik, J. M. McCarthy, "Design of Stephenson linkages that guide a point along a specified trajectory." *Mech. Mach. Theory* 96, 38–51 (2016).

[35] S. S. Baek, F. L. Garcia Bermudez, R. S. Fearing, "Flight control for target seeking by 13 gram ornithopter," 2011 *IEEE/RSJ Int. Conf. on Intell. Robot. Syst.* (IEEE, 2011), pp. 2674–2681.

[36] M. Estrada, E. Hawkes, D. L. Christensen, and M. R. Cutkosky, "Perching and Vertical Climbing: Design of a Multimodal Robot," *IEEE/RSJ Int. Conf. Robot. Autom.*, 2014.

[37] K. Autumn, M. Buehler, M. R. Cutkosky, R. S. Fearing, R. J. Full, D. I. Goldman, R. Groff, W. Provancher, A. A. Rizzi, U. Saranli, A. Saunders, and D. E. Koditschek, "Robotics in scansorial environments," Proc. *SPIE*, vol. 5804, pp. 291–302, 2005.

[38] P. M. Birkmeyer, A. G. Gillies, and R. S. Fearing, "CLASH: Climbing vertical loose cloth," 2011 *IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, pp. 5087–5093, Sep. 2011.

[39] G. Zeglin, "Uniroo--a one legged dynamic hopping robot," B.S. thesis, Mech. Eng., M.I.T., Cambridge MA, 1991.

[40] N. J. Kohut, A. O. Pullin, D. W. Haldane, D. Zarrouk, and R. S. Fearing, "Precise Dynamic Turning of a 10 cm Legged Robot on a Low Friction Surface Using a Tail," *IEEE Int. Conf. Robot. Autom.*, 2013.

[41] D.Rollinson, S.Ford, B.Brown, H.Choset, "Design and modeling of a series elastic element for snake robots", *ASME 2013 Dynamic Syst. Control Conf.* (ASME, 2013), p. v001T08A002.

[42] J. Zhao, T. Zhao, N. Xi, F. Cintron, M. Mutka, L. Xiao, "Controlling aerial maneuvering of a miniature jumping robot using its tail," 2013 *IEEE/RSJ Int. Conf. Intell. Robot. Syst.* (IEEE, 2013), pp. 3802–3807.

[43] N. J. Kohut, D. Zarrouk, K. C. Peterson, R. S. Fearing, "Aerodynamic steering of a 10 cm high-speed running robot," 2013 *IEEE/RSJ Int. Conf. on Intell. Robot. Syst.* (IEEE, 2013), pp. 5593–5599.

[44] M. H. Raibert. *Legged robots that balance*. Cambridge, MA: MIT Press, 1986.

Ultra-High Energy Density TPV Generator for Small Robotic Platforms

Walker R. Chan
Institute for Soldier Nanotechnologies
Massachusetts Institute of Technology
Cambridge, MA
wrchan@mit.edu

Ivan Celanovic
Institute for Soldier Nanotechnologies
Massachusetts Institute of Technology
Cambridge, MA
ivanc@mit.edu

***Abstract*—We present a technology that efficiently harnesses the energy content of hydrocarbon fuels in a volume only a fraction of a cubic inch. A propane-fueled microcombustor heats a photonic crystal emitter to incandescence and the resulting spectrally confined thermal radiation drives low-bandgap PV cells to generate electricity. First, we overcome technical challenges currently limiting thermophotovoltaics: we develop new fabrication processes; we adopt high-temperature alloys to improve thermo-mechanical stability; we adopt commercial polycrystalline tantalum to fabricate large-area photonic crystals; finally, we develop a passivation coating for improved thermo-chemical stability. We demonstrate unprecedented heat-to-electricity efficiencies exceeding 4%, greater than previous 2–3% efficiencies thought to be the practical limit, and we predict that over 12% efficiency is achievable with only engineering optimization. For reference, a 1.5% efficiency corresponds to the energy density of lithium ion batteries. Next, we developed the necessary vacuum packaging to demonstrate the first portable TPV system using a photonic crystal. This work opens new opportunities to free portable electronics, robots, and small drones from the constraints of bulky power sources.**

I. Introduction

In this report, we present our progress towards a mesoscale (length scale on the order of millimeters) generator capable of generating watts of electricity from hydrocarbon fuels, all in a volume smaller than a cubic inch. Mesoscale generators aim to fill the gap between batteries and conventional mechanical generators, where energy demand is burgeoning. Currently, at these small scales, batteries are the only currently available energy sources and already close to their theoretical limit; in contrast hydrocarbon fuels have 60 times the energy density (a 1.5% fuel-to-electricity conversion effciency corresponds to the energy density of lithium ion batteries), thereby promising order-of-magnitude increases in energy storage and conversion. We believe ultra-high energy density power sources are the enabling technology for autonomous mobile microsystems.

Unconventional and fundamentally different energy conversion methods are required at the mesoscale because the vastly different length scales compared to conventional generators result in different geometric scaling (e.g., surface area to volume ratio) and different phenomenological behavior (e.g., laminar and turbulent flow). While micro-mechanical heat engines, fuel cells, and thermoelectrics are actively being applied, in our research, we choose to work within the heat-to-electricity framework of thermophotovoltaic (TPV) energy conversion.

In the TPV energy conversion process, fuel and oxygen react in a microcombustor to generate heat, which brings a photonic crystal emitter to incandescence, and the resulting spectrally confined thermal radiation drives a low bandgap PV cell to generate electricity, as shown in Fig 1a. Our choice of TPV is dictated by the following: TPV has no moving parts and thus does not suffer from increased frictional losses arising from miniaturization; the high-temperature continuous com-

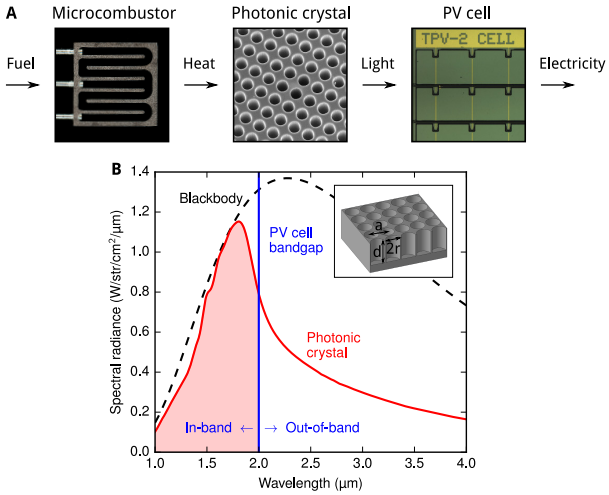


Fig. 1: **(a)** In the TPV energy conversion process, fuel and oxygen react in a microcombustor to generate heat, which brings a photonic crystal emitter to incandescence. The resulting spectrally-confined thermal radiation drives a low bandgap PV cell to generate electricity. **(b)** Simulated spectral radiance at 1000°C normal to the surface of the photonic crystal and a blackbody are plotted. The photonic crystal's emission spectrum is engineered to primarily fall below the bandgap of the PV cell: this in-band radiation, highlighted in green, is converted to electricity. Inset: the photonic crystal is composed of a square array of cylindrical cavities, and its cutoff wavelength can be tuned by varying the radius r, period a, and depth d.

bustion process enables TPV to efficiently process fuel at the mesoscale and be more readily adapted to chemically impure fuel sources such as biofuels; and the physical separation of the thermal and electrical circuits greatly simplifies engineering efforts.

Our overall research objective is address how a practical TPV system would be built. Specifically, our stated research objective is to build a 5-10 W electrical fully functional TPV generator, shown conceptually in Fig. 2, in two steps:

- Generation I of the system will sever external chilled water, vacuum, and fuel/oxidizer supplies. Existing components will be reused wherever possible to save time and money. The end result will be a functional, self-contained, generator but overall system size or weight will not be constrained. Gen I will require a larger heat sink to provide cooling for the cells (200-300 g), vacuum package and housings (300 g), as well as tanks, valves and plumbing (300-400 g).
- Generation II will reduce the overall system weight by incorporating the lessons learned in Gen I, and through increased efficiency to reduce fuel consumption and associated storage overhead, increased electrical output to drive active cooling systems, and tighter system integration and overall optimization to reduce unnecessary weight. Target will be 200 g system (dry) with 150 g of fuel for 24 hour mission at 5 W average power.

These initial objectives proved too ambitious, and work on a fully standalone TPV generator remains ongoing. Our primary accomplishments, listed chronologically below, make great progress towards that objective:

- Developed a detailed model to predict energy density (Wh/kg) of TPV generators by estimating generator mass as well as fuel consumption.
- Developed a fast and low cost method for microcombustor fabrication [1].
- Developed a robust method to integrate the photonic crystal and the microcombustor [1].
- Demonstrated record fuel-to-electricity conversion exceeding 4% in a bench-top experiment [2].
- Developed a propane-air microcombustor [3].
- Developed vacuum packaging [3].
- Designed and fabricated an advanced, filled-cavity photonic crystal capable of higher efficiency [4], [5].

Our revised goal for the MAST program is to demonstrate a nominally portable system with no external connections (e.g. vacuum, chilled water, etc.) capable of producing electricity. Although minimal weight and maximal efficiency will not be our primary objectives, the demonstration will be build such that it can serve as a platform for future work targeting improved efficiency and reduced weight.

II. Related Work

This research builds on previous work at the MIT Institute for Soldier Nanotechnology on high temperature photonic crystals for energy conversion application, high temperature

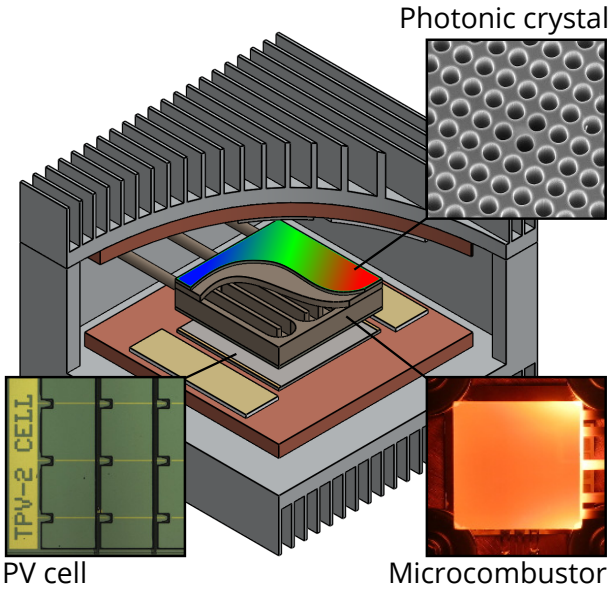


Fig. 2: Rendering of a millimeter-scale TPV generator with integrated microcombustor and photonic crystal hot side and a PV cell.

microcombustor design and fabrication, and TPV system integration.

We have a very strong collaboration with our colleagues from ARL C. Mike Waits and Kimberly Sablon towards demonstrating high-energy density TPV system using liquid fuels, including JP8 military logistics fuel.

III. Approach

With advances in photonic crystals and metamaterials coinciding with advances in low-bandgap III-V PV cells, high efficiency is possible because a large overlap of emission and absorption spectra occurs at readily achievable temperatures of 1000–1200°C, as shown in Figure 1b. Photonic crystals offer near blackbody emission resulting from cavity resonances at the desired wavelengths and near zero emission elsewhere due the low loss of the substrate materials. Their emission spectrum is primarily determined by geometry rather than by specific material properties, allowing the cutoff wavelength to be tuned to match the bandgap of the PV cell. Our 2D photonic crystal, composed of a square array of cylindrical cavities, as shown in the inset of Figure 1b, addresses the challenges of large-area fabrication and integration, good optical performance, and high-temperature stability [6]. In this work, polycrystalline tantalum was used for its machinability, its fast etch rate and good etch selectivity, and the availability of high purity and high metallurgical quality substrates. Our simple fabrication method of interference lithography and deep reactive ion etching enables a uniform, large area (50 mm diameter) to be reliably patterned. To integrate the photonic crystal with the microcombustor, we adopted the industrial process of diffusion brazing. Our photonic crystal is resistant to physical degradation because of its simple geometry, with

a large radius of curvature, and the refractory metal substrate, with a high melting point and low vapor pressure, limiting atomic mobility. Additionally, it is resistant to chemical degradation, such as tantalum carbide formation, because of a conformal hafnium dioxide passivation layer.

Experimental demonstration of the integrated TPV system required design and fabrication of a microcombustor as well as the experimental apparatus, shown in Figure 4. For these experiments, we used pure oxygen in the combustion reaction to emulate exhaust recuperation, which is required in a portable, air-breathing system. Propane and oxygen entered the microcombustor through the inlet tubes, flowed through an internal serpentine channel where they reacted on the catalyst-coated walls (5% platinum on porous alumina), and exited though the outlet tube [7]. Heat was conducted through the channel walls to the photonic crystal emitters bonded to the top and bottom surfaces of the microcombustor, as shown in Figure 4. The photonic crystal emitted spectrally confined thermal radiation matched to the InGaAs cells ($E_g = 2.0$ μm) mounted below the assembly [8]. The experiment was housed in a custom vacuum chamber to suppress convection and prevent degradation of the photonic crystal by reaction with air. A window was used in place of the top set of cells.

A. *Design and Simulations*

The microcombustor reacts propane and oxygen to heat the photonic crystal to ∼1000°C. We used a catalytic design because of the difficulty maintaining a flame at the mesoscale due to thermal and radical quenching at the walls. A planar serpentine channel with catalyst-coated walls provided sufficient interaction time for complete combustion while fitting within external dimensions, which were matched to those of the available PV cells. The channels were dimensioned to provide sufficient length such that the residence time of the gases was much greater than the time for gases to diffuse across the channel, as described in Ref. 7. The microcombustor was suspended by three tubes that doubled as fluidic connections to minimize conductive heat losses. To improve temperature uniformity, a symmetric design was used with two inlets and a central outlet: the increased heat production near the inlet tubes compensated for the increased heat loss due to the edges. In the inlet tubes, propane was delivered via a fine capillary tube run inside the outer tube, and oxygen was delivered via the annulus formed between the capillary and outer tube. This tube-in-tube configuration was necessary to prevent flashback and premature combustion in the inlet tubes.

We modeled the microcombustor with a custom heat transfer code incorporating the radiation from front and back emitters and the edges of the microcombustor, conduction through the support tubes, and heat carried out by the hot exhaust gases. The hemispherically-averaged emissivity for the photonic crystal structure was computed by the Fourier Modal Method, using a freely available software package, in which the optical dispersion was captured with a Lorentz-Drude model fitted to unstructured tantalum. Simulated and measured normal-incidence emissivities, plotted in Figure 5b, agree well.

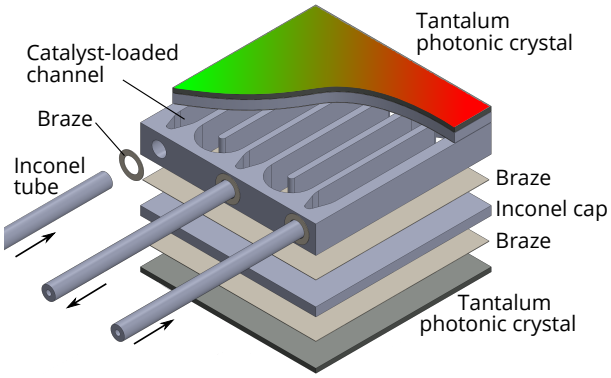


Fig. 3: Exploded view of integrated microcombustor and photonic crystal.

Then, we used ray optics to accurately incorporate multiple scattering effects between the emitter and PV cell, assuming purely diffuse emission and reflection. The PV cell was modeled using a single diode equivalent circuit methodology [9]. The microcombustor temperature, assumed to be uniform, was solved self-consistently. Simulated electrical power output and temperatures are shown in Figure 5a; the shaded regions indicate possible ranges of experimental parameters, primarily the emissivity of the edges of the microcombustor.

B. *Fabrication*

Diffusion brazing is an industrial process similar to brazing, however addition and subsequent diffusion of melting point depressants out of the brazing alloy increases its remelt temperature—allowing the assembly to be reliably operated *above* the original brazing temperature. In this work, we chose BNi-2 (Lucas-Milhaupt) for the brazing alloy, with a solidus temperature of 971°C and a liquidus temperature of 999°C, and the following composition: 7% chromium, 3% boron, 4.5% silicon, 3.0% iron, and balance nickel. The braze alloy is subjected to a prolonged anneal above its liquidus, during which the silicon and boron diffuse out and the molten alloy undergoes isothermal solidification. Assuming the silicon and boron are completely removed, the remelt temperature exceeds 1400°C. The increase in remelt temperature enables us to use the same braze alloy for multiple brazing steps, to avoid exposing the photonic crystal to a higher temperature than absolutely required, and to perform the brazing in a low-cost furnace which was limited to 1200°C even through the target operating temperature is 1000–1200°C.

We were concerned that significant stresses, and therefore deflections, could occur in the final brazed assembly owing to the differential thermal expansion between Inconel and tantalum. Inconel was required for the microcombustor for its high-temperature oxidation resistance, low cost, and machinability; tantalum was required for the photonic crystal for its low vapor pressure, low optical loss, and etchability. We modeled the deflection of an Inconel-tantalum bimetallic strip in SolidWorks Simulation using available material properties and found significant deflection. To test the validity of the

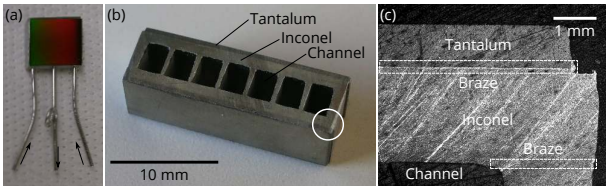


Fig. 4: (a) Completed hot side with a diffraction pattern visible on the photonic crystal. Arrows indicate direction of gas flow. (b) Cross section with the channels visible. (c) A micrograph of the indicated corner of the hot side assembly with the tantalum-Inconel braze joint visible and the Inconel-Inconel braze indicated but not visible.

simulation, two test coupons with 1.02 mm and 3.18 mm thick Inconel and 0.51 mm thick tantalum were prepared, and deflections of 1.1 and 0.2 mm were measured, 2–3 times greater than simulated values possibly because of inconsistent material properties or plastic deformation. To prevent deflection, we balanced the stress with a symmetric design and eliminated long spans by brazing each channel wall to both the top and bottom Inconel caps, as shown in Fig. 1(b).

The microcombustor channels were designed as described in Ref. 7 and fabricated by abrasive water jet cutting from Inconel sheet stock. The holes for the tubes were reamed over-sized to ensure a consistent 25 μm gap between the tube and hole so that the braze would reliably flow by capillary action. Tubing was commercially available and was simply cut to length. We bent a loop in the center tube to relieve stress arising from differing thermal expansion between inlet and outlet tubes, visible in Fig. 4(a). The photonic crystal was fabricated separately, as described in Ref. 6. Braze preforms were fabricated from foil by photochemical machining using dry film photoresist and ferric chloride etching solution commonly used for printed circuit board fabrication. Preforms were dimensioned to deliver a slight excess of braze alloy to the joint.

Brazing was conducted in three steps: tubes were brazed to the channels, caps were brazed to seal the channels, and the photonic crystals were brazed to the completed microcombustor. Jigs were used to hold components in place for each of the steps. The jigs used for the first and second brazing operations were machined from Inconel. The one for the final brazing operation was machined from tantalum to avoid contamination of the photonic crystal, as some outgasing of the Inconel was observed.

Brazing was performed a quartz tube furnace evacuated by a turbomolecular pump. We relied on the high temperature and high vacuum to shift the chemical equilibrium to favor the dissociation of surface oxides before the braze alloy melted. flux and reactive atmospheres (e.g. hydrogen) were avoided to prevent contamination of the photonic crystal. After pump-down, the furnace was ramped at 10°C/minute, with one hour stops at 350°C and 500°C for degassing, to a final brazing temperature of 1100°C. When the brazing temperature was reached, furnace pressure would initially

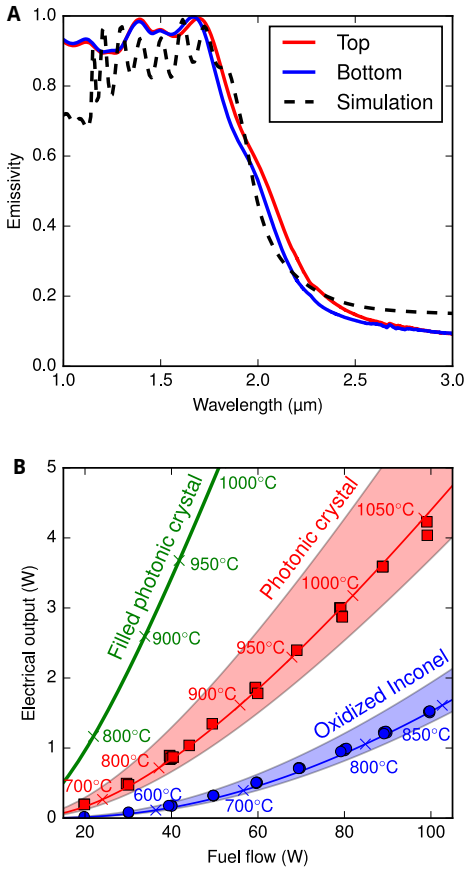


Fig. 5: (a) Measured and simulated room temperature, normal incidence emissivity of the photonic crystal. (b) Measured and simulated electrical power output as a function of fuel flow. Points are experimental results. Lines and shaded bands are simulation results, with the bands indicating a range of uncertain parameters. Simulated temperatures are indicated. Note that the filled photonic crystal (green line) has an electrical power output of 12.6 W at 100 W of fuel flow (not shown).

spike to $\sim 5 \times 10^{-5}$ Torr then reduce to $\sim 3 \times 10^{-6}$ Torr. The temperature was held at 1100°C for two hours to ensure full diffusion before returning to room temperature.

During process development, we fabricated ten hot side assemblies with bare tantalum substituted for the photonic crystal, and experienced high fabrication yield as defined by visual inspection and helium leak detection. A cross section of one is shown in Fig. 4(b) and (c). One microcombustor without a photonic crystal was operated for one week without failure or visual damage. Once we had perfected the brazing process, we used a photonic crystal instead of bare tantalum. The finished assembly is shown in Fig. 4(a). Reflectance measurements of the photonic crystal before and after brazing confirmed no degradation of optical properties of the photonic crystal.

In summary, we adopted the industrial process of diffusion brazing and fabricated an integrated Inconel microcombustor

and tantalum photonic crystal to serve as the hot side of a millimeter-scale TPV generator. In diffusion brazing, fast-diffusing elements contained in the alloy diffuse out of the joint during heating, thus increasing the remelt temperature of the braze above the original brazing temperature. This approach represents a breakthrough for TPV because it integrates the microcombustor and photonic crystal, and is fast, simple, and reliable. Diffusion brazed Inconel is applicable to microcombustor fabrication for other applications, and to other millimeter-scale devices operating at high temperatures.

IV. Results

To ignite it, the microcombustor was heated to approximately 400°C with a halogen lamp through the window. Above that temperature, the propane kinetics over the catalyst were sufficient for autothermal operation, and the halogen lamp was shut off. Propane flows, corresponding to a total latent heat input of 20–100 W, were increased in small increments while maintaining an oxygen flow of 7.5 times that of propane (an equivalence ratio of $\phi = 1.5$), and steady-state electrical output at the maximum power point was measured. We characterized both a microcombustor with and without a photonic crystal emitter. For the microcombustor without the photonic crystal, the bare Inconel surface was oxidized by operation in air until visibly black (emissivity of $\varepsilon \approx 0.8$) and used as the emitter. Electrical measurements are shown in Figure 5b, scaled for a full set of cells. We measured 4.3% fuel-to-electricity conversion efficiency with the photonic crystal emitter and 1.5% with the oxidized Inconel emitter, for a fuel input of 100 W.

We used the model to study heat flows within the TPV systems with the photonic crystal and oxidized Inconel emitters, shown in Table I. Surprisingly, in-band radiation increased nearly three times in the case of the photonic crystal even though its hemispherically- and wavelength-averaged in-band emissivity was lower ($\varepsilon = 0.59$ compared to $\varepsilon = 0.8$). Because the photonic crystal suppressed out-of-band radiation, the temperature increased by approximately 200°C, resulting in increased thermal radiation. From a practical point of view, the low out-of-band emissivity allows for simpler system design without the need for photon recycling via a cold side filter placed between the emitter and PV cell, a difficult approach because of the wide range of angles and wavelengths that need to be filtered with high selectivity and low loss. Moreover, because of the photonic crystal's high in-band emissivity, electricity generation can compete favorably with other heat loss mechanisms and a readily achievable emitter temperature of 1000°C should provide greater than 500 mW/cm^2 of cell area output. These factors are important in a portable system: the low out-of-band emissivity minimizes the waste heat generated by the PV cell and thus the heat sink size; and the high in-band emissivity minimizes the active area required for a specified electrical output as well as increasing efficiency.

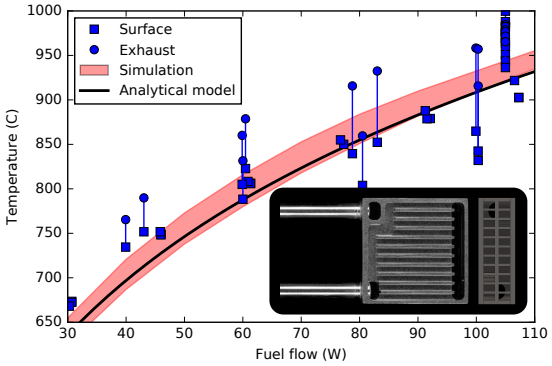


Fig. 6: Simulated and measured operating temperature for the propane-air microcombustor as a function of fuel flow (calculated from the lower heating value). Inset: microcombustor structure.

A. Propane-air microcombustor

The propane-oxygen microcombustor was modified to operate on propane-air. Specifically, a two layer design was required in order to attain sufficiently long residence time and sufficiently short diffusion time to react 100 W of propane flow with air.

Two simulations were used. First, we used a simple analytical heat balance model, which assumed a uniform temperature across the burner surface. Second, we developed a reduced-order modeling tool that enables the evaluation of different reactant compositions and geometry modifications on system thermal efficiency and surface temperature uniformity. These results are shown in Fig. 6, where the shaded region represents the range of surface temperatures predicted by the second model.

The microcombustor was tested in the vacuum chamber described in the next section. A thermocouple was spot welded to the surface to measure temperature and another thermocouple was inserted into the exhaust tube during some experiments to measure exhaust gas temperature. The microcombustor was ignited by bubbling the air through methanol, which reacted at room temperature over platinum, until heated to around 300°C, at which point propane-air operation was possible and the methanol bubbler was bypassed. Exhaust and surface temperatures for a range of propane flows are shown in Fig. 6.

B. Vacuum Packaging

1) Conflat chamber: In order to prevent degradation of the photonic crystal emitter and to suppress convective losses, vacuum packaging is required. To test the feasibility of vacuum packaging we assembled a microcombustor and hot filament ion gauge in a Conflat (CF) tee, as shown in the inset of Fig. 7. A zirconium based getter (Saes) was also placed in the chamber. A soft copper tube connected the test chamber to a turbomolecular pump and residual gas analyzer. With the turbomolecular pump running, the chamber was heated to 350°C to degas and to activate the getter. In addition, the

Emitter	Oxidized Inconel	Photonic crystal	Filled photonic crystal
Emitter emissivity ($\varepsilon_{in}/\varepsilon_{out}$)	0.80/0.80	0.58/0.18	0.92/0.16
Edge emissivity	0.80	0.55	0.15
Fuel	100 W	100 W	100 W
Exhaust	15.2 W	19.7 W	22.2 W
Conduction	3.6 W	4.6 W	5.1 W
Side radiation	27.3 W	32.8 W	12.5 W
View factor loss	4.8 W	3.6 W	5.4 W
Out-of-band radiation	46.9 W	24.5 W	15.2 W
Cell inefficiencies	4.3 W	10.6 W	27.1 W
Electricity	1.5 W	4.4 W	12.6 W
Fuel-to-electricity efficiency	1.5%	4.4%	12.6%
Temperature	843°C	1054°C	1169°C

TABLE I: The simulated systems with the oxidized Inconel emitter, photonic crystal emitter, and filled photonic crystal emitter. The hemispherically averaged in-band and out-of-band emissivities of the emitters, edge emissivity, power distribution, efficiency, and temperature are listed.

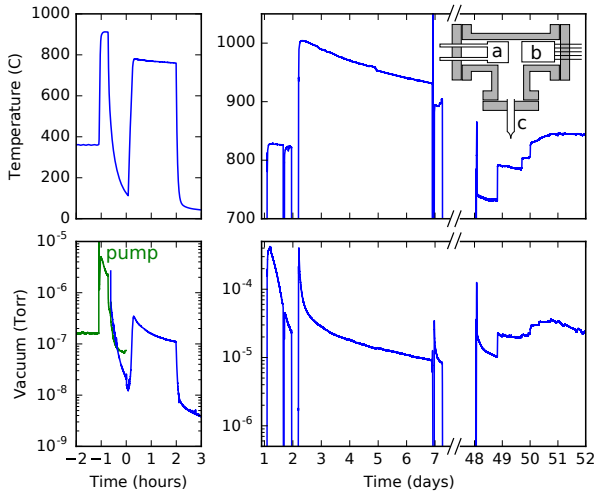


Fig. 7: Temperature and vacuum during a 50+ day experiment, with time measured relative to pinch-off ($t = 0$). Inset: the experimental setup was composed of a microcombustor (a), ion gauge (b), and copper pinch-off (c).

burner was ignited and the ion gauge was electrically heated. At the beginning of the bakeout, water vapor dominated; at the end, hydrogen and carbon monoxide dominated. The hydrogen was determined to be permeating through the walls of the microcombustor by injecting pulses of deuterium into the combustion reaction and observing the isotope ratio in the vacuum chamber. At the end of the bakeout, the copper tube was pinched off with a hydraulic crimper, isolating the vacuum chamber from the pump. Temperature and vacuum near pinch-off are shown in Fig. 7.

After pinch-off the microcombustor was run for six days without degradation of the vacuum. In fact, an improvement in vacuum was observed because the getter continued to act as an internal pump. The apparatus was left for nearly 40 days, during which the vacuum level was $\sim 10^{-8}$ Torr. The microcombustor was successfully reignited and run without

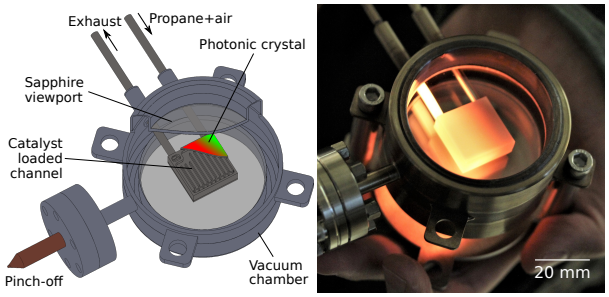


Fig. 8: Cutaway rendering and photograph of the vacuum packaged microcombustor with integrated photonic crystals in operation.

vacuum degradation, and stepped through several fuel flows.

2) Windowed chamber: After validating the microcombustor and vacuum packaging strategy, we designed and fabricated an integrated unit with a photonic crystal emitter and sapphire windows to transmit the infrared radiation to PV cells. The package, shown in Fig. 8, was fabricated from commercial sapphire viewports and custom components by welding and brazing. Although vacuum level could not be directly measured, no discoloration of the microcombustor or photonic crystal (an indication of oxidation) was observed.

V. Conclusion

The primary research accomplishments described in the report are the following:

- Developed a diffusion brazing process for microcombustor fabrication and integration of the photonic crystal and microcombustor.
- Demonstrated record fuel-to-electricity conversion exceeding 4% in a bench-top experiment.
- Demonstrated robust vacuum packaging of integrated microcombustor and photonic crystal.

The next step is to achieve a portable demonstration are the following:

- Purchase and assemble COTS miniature plumbing and balance of plant components.
- Design and fabricate a recuperator for improved propane-air efficiency.

We are in the process of exploring the potential for technology transition. Our target applications are providing power for dismounted soldiers, and for military robotics and other specialized equipment. A soldiers mission can require that he carry all the supplies he needs for 72 hours on his back, including 10-20 lbs of batteries. Not only does this burden the soldier, reducing his mobility and increasing his chance of injury, it prevents him from carrying other equipment, food, water, and ammunition. The power demands will continue to grow as the Army grows more connected (radio communications are largest electrical load) and fields more robotic platforms, drones, and other battery-powered equipment. A small, portable TPV generator could be used to trickle charge the soldiers batteries on the move. By replacing many of the spare batteries with fuel, which has a greater energy density, we predict that we can reduce the total weight of the power sources (batteries, generator, and fuel) by 75%. To power a 10 W load (average power load of a soldiers equipment) would require a generator weighing 250 g and 350 g of fuel per day.

REFERENCES

[1] W. R. Chan, V. Stelmakh, W. R. Allmon, C. M. Waits, M. Soljacic, J. D. Joannopoulos, and I. Celanovic, "An integrated microcombustor and photonic crystal emitter for thermophotovoltaics," in *Journal of Physics: Conference Series*, vol. 773, p. 012108, 2016.

[2] W. R. Chan, V. Stelmakh, M. Ghebrebrhan, M. Soljačić, J. D. Joannopoulos, and I. Čelanović, "Enabling efficient heat-to-electricity generation at the mesoscale," *Energy Environ. Sci.*, vol. 10, no. 6, pp. 1367–1371, 2017.

[3] W. R. Chan, V. Stelmakh, S. Karnani, C. M. Waits, M. Soljacic, J. D. Joannopoulos, and I. Celanovic, "Towards a portable mesoscale thermophotovoltaic generator," in *Journal of Physics: Conference Series*, 2017 (submitted).

[4] Y. X. Yeng, J. B. Chou, V. Rinnerbauer, Y. Shen, S.-G. Kim, J. D. Joannopoulos, M. Soljacic, I. Čelanović, M. Soljačić, and I. Čelanović, "Omnidirectional wavelength selective emitters/absorbers based on dielectric-filled anti-reflection coated two-dimensional metallic photonic crystals," in *Proc. SPIE*, vol. 9170, p. 91700X, 2014.

[5] V. Stelmakh, W. R. Chan, M. Ghebrebrhan, M. Soljacic, J. D. Joannopoulos, and I. Celanovic, "Fabrication of an Omnidirectional 2D Photonic Crystal Emitter for Thermophotovoltaics," in *Journal of Physics: Conference Series*, vol. 773, pp. 3–6, 2016.

[6] V. Stelmakh, V. Rinnerbauer, R. D. Geil, P. R. Aimone, J. J. Senkevich, J. D. Joannopoulos, M. Soljacic, I. Celanovic, M. Soljačić, and I. Celanovic, "High-temperature tantalum tungsten alloy photonic crystals: Stability, optical properties, and fabrication," *Applied Physics Letters*, vol. 103, no. 12, pp. 0–4, 2013.

[7] W. R. Chan, B. A. Wilhite, J. J. Senkevich, M. Soljacic, J. Joannopoulos, and I. Celanovic, "An all-metallic microburner for a millimeter-scale thermophotovoltaic generator," in *Journal of Physics: Conference Series*, vol. 476, p. 12017, 2013.

[8] D. Wilt, R. Wehrer, M. Palmisiano, M. Wanlass, and C. Murray, "Monolithic interconnected modules (MIMs) for thermophotovoltaic energy conversion," *Semiconductor Science and Technology*, vol. 18, no. 5, pp. S209–S215, 2003.

[9] W. R. Chan, R. K. Huang, C. A. Wang, J. Kassakian, J. D. Joannopoulos, and I. Celanovic, "Modeling low-bandgap thermophotovoltaic diodes for high-efficiency portable power generators," *Solar Energy Materials and Solar Cells*, vol. 94, no. 3, pp. 509–514, 2010.

Integrated Air-Surface Operation for Micro Air Vehicles

MCE-16-17.4.4

Mark. R. Cutkosky, Morgan T. Pope, Hao Jiang,
Matthew A. Estrada, William R. T. Roderick
Department of Mechanical Engineering
Stanford University
Stanford, CA, USA
cutkosky@stanford.edu

J. Sean Humbert
Department of Mechanical Engineering
University of Colorado Boulder
Boulder, CO, USA
sean.humbert@colorado.edu

***Abstract*— We present adhesion models and technologies for small fliers that land, perch and crawl on outdoor surfaces using arrays of micro-spines. Statistical models of surfaces and spine/surface attachment provide insights for perching mechanism design and for the control of fliers just prior to landing. The design of mechanisms and maneuvers also draws insight from biological sources including birds and insects. Dynamic models of the landing and attachment process lead to a computed "envelope" of initial conditions for which perching is likely to succeed. Onboard sensors provide sufficient information to detect attachment failure or success, and to trigger transitions among states including approach, landing, perching, crawling and take-off.**

***Index Terms*—MAV, Perching, Climbing**

INTRODUCTION

Perching enables micro air vehicles (MAVs) to conserve energy or recharge their batteries while performing useful tasks, such as sensing and communication [1-13]. In disasters, standing walls – free from debris because they are vertical – are typically present and suitable for perching. In nature, all flying animals also have a ground locomotion strategy, and a surprising number of climbing animals, from ants to marsupials to snakes, have an aerial locomotion strategy [14]. The strategies animals use for successfully navigating the transitions between aerial phases and surfaces provide insights for the mechanisms we design and the maneuvers we perform.

We aim to increase the versatility of landing, perching and deploying devices on flat/curved surfaces and objects including walls, pipes, and power lines; our approach increases the robustness of perching using on-board sensing and computation. This involves three main tasks: (i) modeling and analysis of adhesive technology; (ii) modeling hybrid dynamics to go beyond aerodynamic control using inertial/mechanical methods; (iii) employing new sensing strategies (e.g. vision, accelerometer arrays) to operate robustly without 3D motion capture systems. In this final report, we focus on the first two tasks and expand our results to a range of perching vehicles, maneuvers, and strategies.

As a starting point, we study rough surface properties to predict micro-spine gripper adhesion capabilities and give insights on mechanism designs. Thin film gecko adhesive grippers are also developed to enable attachment to smooth cylindrical surfaces. We integrate these grippers onto quadrotor platforms to perch and climb on various surfaces, both indoors with motion capture systems and outdoors with only onboard sensing. Perching mechanisms with inertial control are developed for more precise surface alignment inspired by gliding geckos.

RELATED WORK

The work takes advantage of ongoing MAST research by Cutkosky (MCE 15 4.4), Kumar (MCE 4.5) and Humbert and Yoon (Optic Flow Motion Sensing and Processing, SPP 16-1.2), Najafi (HAIR Sensors for Inertial Navigation, SPP 16-1.1), Humbert and Chopra (MAST Scale Platform Flight Dynamics and System ID, MCE 16-1.5), and smart skin (arrayed accelerometer) research presented in a separate white paper by Humbert and Chang (Stanford Aero/Astro). New modeling, prototype fabrication and testing is performed in close collaboration between the Stanford Biomimetics and Dexterous Manipulation Lab, the UMD Autonomous Vehicle Laboratory, and the UPenn GRASP Lab. This work also leverages interactions with NASA/JPL (Dr. Aaron Parness) to use opposed directional adhesives for grappling debris in orbit. The same energy absorption and gripping principles (at very different physical scales) are used for UAVs and for space adhesion.

We also continue an informal collaboration with Robert Full's Lab at UC Berkeley, where we build on their insights into

animal locomotion (e.g., gliding geckos [15], to create robots that validate dynamic hypotheses).

APPROACH

Micro-spine Adhesion Modeling

Micro-spines can attach to bumps and pits (asperities) on rough surfaces as the tips travel along the surfaces. However, the process is inherently stochastic [16-18]. Thus, the asperity distributions need to be studied to inform the adhesion capability of a group of spines.

With the assumption that asperities are randomly located on surfaces, we can model the asperity occurrence as a spatial Poisson process with a mean interval distance. Thus, the interval distance between adjacent asperities can be modeled as an exponential distribution. Moreover, many asperities have non-zero lengths such that a micro-spine can catch on any point along an asperity. Thus, we can model the asperity distribution as a 2-phase phase-type distribution with a parameter α to balance the probability between how far the asperities are spaced apart (exponential term) and how long an asperity can be (special term).

$$\begin{aligned} f(\alpha,\lambda;x) &= \alpha f_exp + (1-\alpha)\, f_special \\ f_exp &= \lambda e\text{^}(-\lambda x) \\ f_{special} &= \delta(x) \end{aligned} \tag{1}$$

where x is a random variable describing the interval distance between adjacent asperities; λ is the average number of asperities per unit length on a surface (e.g. 1mm); 1 - α is the probability of immediate engagement when a micro-spine arbitrarily lands on a surface; and δ is the impulse function which can be treated as an extreme case of exponential distribution with λ equal to infinity.

There are two failure cases for spine/asperity contact: (1) the spine holds onto an asperity until the contact force exceeds the asperity strength (asperity failure) and (2) the spine itself breaks or slips due to its elastic/plastic bending before the asperity fails (spine failure). The asperity strength distribution is not accurately described by a standard distribution function, nor can it easily be derived from physical observations. Thus, we use raw measurement data for the asperity strength distribution. For simplicity in many cases it is feasible to use a uniform distribution as an approximation. The adhesive stress of a group of directional spines can then be calculated based on the asperity spatial and strength distributions:

$$E[\sigma] = \frac{\int_0^{d_{stroke}} f(\alpha,\lambda;x)\, F(x)\, Pr(y \geq F(x))\, dx}{(d_{unit}+d_{flexure}+d_{stroke})\, w} \tag{2}$$

$$F(x) = K\,(d_{stroke} - x) \tag{3}$$

where $E[\sigma]$ is the expected adhesive stress at full travel; $d_{flexure}$ is the flexure length of each spine unit; d_{stroke} is the maximum travel of the tile base, as well as the interval spacing between spine rows; $f(\alpha,\lambda;x)$ is the probability density of engagement with travel x; $F(x)$ is the adhesive force of a single spine free traveling x to engage and being loaded to the maximum traveling distance; $Pr(y \geq F(x))$ is the cumulative probability of an asperity stronger than $F(x)$; w is the width/thickness of a single micro-spine; and K is the spine stiffness. Similarly, the standard deviation of the adhesive stress can be calculated based on the two asperity distributions.

In a more complex recursive model, we also consider the effect of a spine re-engaging on another asperity after it fails on the previous one [19]. The entire loading process along the stroke can be described as a Markov Decision Process (MDP). The computational complexity of the MDP grows exponentially along with finer discretization of the whole stroke. For example, where a simple single-layer model takes approximately one second to compute, a model with four layers takes over an hour. Therefore we use the single layer model on most surfaces, although a 3-layer model is better for weak surfaces (e.g. adobe).

Perching and Climbing with Quadrotor Platforms

For quadrotor perching outdoors, we have used opposed-spine mechanisms for vertical and overhanging surfaces. For combine perching and climbing, we focus on unidirectional spines and tune the lift and stroke timing to achieve reliable attachment [13].

The opposed-spine mechanisms consist of two micro-spine tiles constrained on a linear rail (Fig. 1). Reset springs pull the spine tiles towards their starting position. The actuating truss is bi-stable. In the neutral state, it is pulled by the reset springs towards the horizontal hard-stops. This is the first stable position. There is a pre-stretched preload spring in line with the tendon connecting each tile to one end of the truss. When the truss collapses, the two ends of the truss slide apart for 1cm to pull the tiles inward for preload. The truss can fully collapse and pass the horizontal position. After that the two tendons tend to pull the truss such that the apex remains in the over-center position. A vertical hard-stop holds the mechanism in place. This is the second stable position. Load is directly exerted on the frame of the mechanism through a loading tendon that is routed around the truss for 2 cycles. This eliminates slack, which is otherwise caused by the unequal travels of the apex of the truss and the two ends of the truss. During release, a tiny force is needed to pull the apex of the truss up so that the reset spring pulls it back to its first stable position. At this point the spine tiles slide apart to detach from the engaged asperities. The release force can be exerted by a small servo motor and cam or an SMA actuator. We have demonstrated this mechanism integrated with off-the-shelf lightweight quadrotors that perch on vertical and inverted rough surfaces outdoors.[1]

[1] https://www.youtube.com/watch?v=7cKlbmfY6cQ

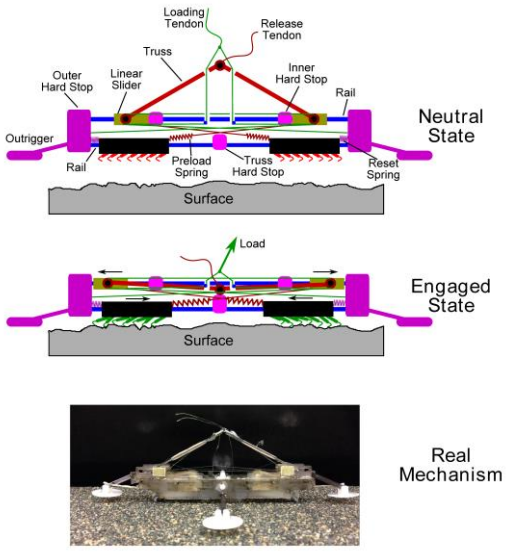


Fig. 1. Diagrams and a photo of an opposed micro-spine perching mechanism.

For combining perching and climbing, particularly on vertical surfaces, we use a slightly different strategy than perching on inverted surfaces, which requires strong normal adhesion. Instead, we equip a quadrotor with unidirectional spines for both perching and climbing. The demonstration platform, Stanford Climbing and Aerial Maneuvering Platform (SCAMP), is shown in Fig. 2 [13]. Rotor thrust enables SCAMP to climb on a wider range of surfaces than most climbing robots, as it increases the number of usable asperities. Based on the coefficient of friction between a spine and the local surface, as well as the loading angle (determined in normal climbing applications by tail length and center of mass location), some segments of the wall are at an angle shallow enough to support a climbing grip. By adding rotor thrust, we can change the effective loading angle and turn some previously unusable asperities into usable ones. Modeling this effect predicts a nonlinear increase in asperity density with small amounts of thrust on certain surfaces. Moreover, using the onboard accelerometer, we are able to detect when a step is missed and a fall begins. Turning the rotors on for a short time upon detection can help SCAMP return to contact with the surface, recover from failure, and enable reattachment for continued climbing.

Curved Surface Perching and Gripping

We additionally have studied curved surface perching in both indoor and outdoor environments. Adapting a previously developed gecko adhesive curved surface gripper [20], we use quadrotors to demonstrate perching on target pipes, approaching from underneath, using a motion capture system (via collaboration with UPenn). In other work we have developed grippers for free-flying vehicles that will dock on target objects in free floating environments via collaboration with NASA. Based on the adhesion model of the gripper and the vehicle dynamics, a landing envelope was characterized and compared with experiments. In addition to the regular energy absorption and mechanism engagement concerns discussed above, the initial conditions for successful landing are also dependent on the direction of the angular velocity of the vehicle and the translational misalignment between the vehicle and the target. A "rolling-in" behavior is much more desirable than a "rolling-out" approach, thus providing insights for some last-minute control to compensate for different rolling behaviors [21]. In one more application, a micro-spine based curved surface gripper was developed to enable attachment to tree trunks and power poles. Combining the insights of the flat surface micro-spine grippers and the study of avian feet, we integrated a micro-spine tile module on the end of each finger such that the finger takes care of surface conformation. The micro-spines are then loaded in the same way as for flat grippers. As the curvature becomes more pronounced a larger surface normal preload is desired to increase the number of usable asperities in a similar way as for the aforementioned rotor-assisted climbing. Although an increased normal preload tends to push the gripper fingers away from the cylindrical surface, the increased density in usable asperities compensates for this effect and provides a larger adhesion [22].

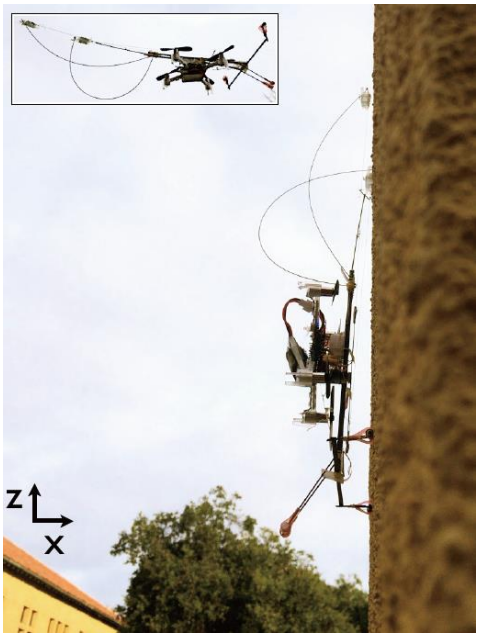


Fig. 2. The Stanford Climbing and Aerial Maneuvering Platform perched on the Stanford clock tower.

Perching Maneuver Correction with Inertial Control

The perching process, especially considering the impact, is often in the order of 50ms. With most actuators and controllers, it is difficult to correct any linear/angular position/velocity error in this duration. However, errors can possibly be compensated in advance of the perching impact to minimize the chance of failure and after the impact during the energy absorption process. We studied the perching behaviors of flying geckos who utilize their tails to balance their body during the perching process. Similarly, we implemented inertial reorientation on a ballistic percher in collaboration with UC Berkeley's PolyPEDAL Lab. In the air and at first contact, without much drag or other actuation, the angular momentum of the system is conserved. Thus, by controlling the motion of the inertial tail, we can correct the angular misalignment of the percher and stabilize its attitude. The tail control can be continued after the perching to further reduce the rebound of the percher.

RESULTS

Spine Adhesion Modeling

The shear adhesion of a group of 20 spines was measured to further verify the stochastic adhesion model. Five micro-spine tiles with different parameters were made using the Shape Deposition Manufacturing (SDM) [23,24] process. They all share the same unit length (8mm), maximum spine force (3N) at full stroke, and the total number of spines (20). Each tile has a unique stroke ranging from 1.5mm to 8.5mm, and thus different stiffness. During tests, each tile is placed arbitrarily on 3 different surfaces (a concrete block, a piece of roofing shingle, and a piece of 80-grit sand paper) and loaded in shear to full travel/stroke. The shear adhesive stresses were recorded and compared to the simplified model, as shown in Fig. 3. The discrete points and error bars are the mean and 1 standard deviation of the measured adhesive stress. Each data point corresponds to 20 test trials.

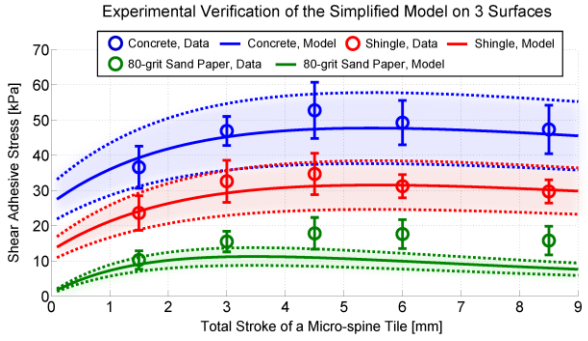


Fig. 3. Experimental verification of the simplified model. Each data point corresponds to a unique tile parameter and 20 test trials. The solid and dotted lines are the predicted mean and 1 standard deviation of the adhesive stress with the model. The blue points and lines are with a fine concrete block, the red ones are with a piece of roofing shingle, and the green ones are with a piece of sand paper.

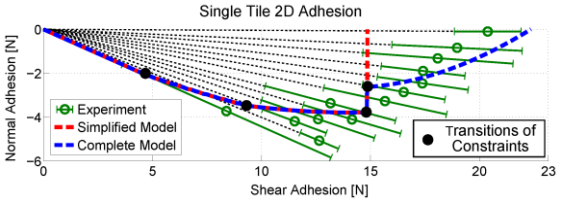


Fig. 4. Experimental verification of the 2D adhesion model.

Only for weak surfaces like sandpaper is the single-layer model sufficiently inaccurate to warrant the multi-layer model.

Beyond shear stress characterization, the 2D adhesion performance of both unidirectional spines and opposed-spine mechanisms were modeled and measured. A group of 20 spines with the same parameters as in the previous experiments were dragged on a concrete surface in different loading directions. An ATI Gamma SI-130-10 force sensor beneath the surface measured the external load. One hundred data points were taken for different loading angles and averaged within every 2° angular bin. A comparison of predicted 2D adhesion between the model and the experiment is shown in Fig. 4. The error bars indicate 1 standard deviation along the force vectors. The error is large when the tile loading angle is close to the loading angle limit of the surface, when small perturbations can cause spines to slip.

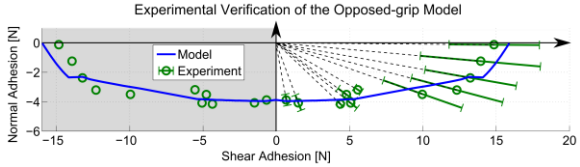


Fig. 5. Experimental verification of the opposed spine mechanism adhesion model.

Additional experiments are conducted with two opposed-spine tiles, again with 20 spines on each tile and the same parameters as before. One hundred points are recorded with various loading orientations from pure shear to pure normal. With the assumption that the surface is isotropic, and the two sides of the spine tiles are almost identical (achievable with good manufacturing control), we only take data towards one side and mirror the result for the other side. The comparison between the model and the experiments are shown in Fig. 5. Each data point is the average of the actual data within every 5° angular range The error bars indicate 1 standard deviation in the corresponding loading angle. Note that in this case, the stochastic variation is least when loading in the normal direction, where small angular perturbations have little effect.

Natural Curved Surface Perching

We investigated the design and modeling of a claw or spine-based gripper for perching on rough, curved surfaces. Drawing inspiration from the opposed grip techniques found in animals, we focused on design considerations associated with surface geometry and preload. A model elucidates the relationship between these variables, and a mechanism demonstrates the effectiveness of the opposed grip technique. Fig. 10 shows the opposed gripper prototype along with typical performance on several bark surfaces.

Perching on rough, vertical curved surfaces depends strongly on multiple interconnected factors including contact geometry in relation to platform size, as well as preload. On relatively flat surfaces, increasing the shear preload can substantially increase the maximum outward normal load. As the surface curvature increases, normal preload dominates and facilitates a firm grip. In all cases, it is desirable to provide the spines with a long suspension travel so that they can reach deep into features for maximum loads. Future work will focus on improving claw/surface interaction models, with extensions to multiaxial loading and improved claw and spine designs for arboreal fliers.

Perching and Climbing

Vertical and inverted surface perching experiments are conducted with a 150g RC controlled quadrotor with a 15g opposed micro-spine mechanism on the top. There are 8 spines on each side. There is a group of single-sided spines on a tail that is deployed at the front end of the quadrotor. The quadrotor perches on both vertical and inverted rough surfaces quasi-statically. For vertical surface perching, upon a gentle impact with the surface using the front tail, the tail spines quickly catch on asperities such that the quadrotor starts to pivot forward about the tail contact point to align with the surface. The thrust pushes the apex of the truss to make the opposed spines engage with the surface. After the thrust is turned off, the tail spines continue to provide shear adhesion, and the opposed spine mechanism mainly provides normal adhesion to counterbalance the pitch-back moment. A typical flight and perching sequence is shown in Fig. 6. During release, a servo on the mechanism opens the truss and detaches the mechanism from the surface. The thrust cannot

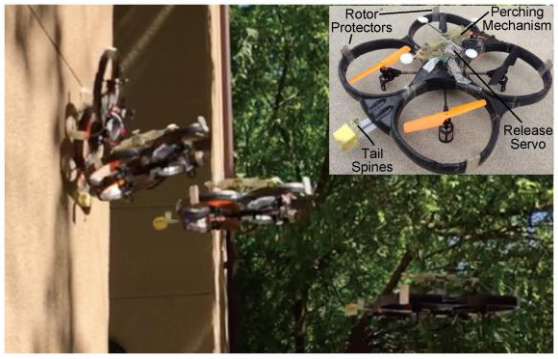


Fig. 6. The flight sequence of a quadrotor perching on a vertical rough surface. The quadrotor pivots forward on its tail spines to align with the surface and trigger the truss collapsing mechanism.

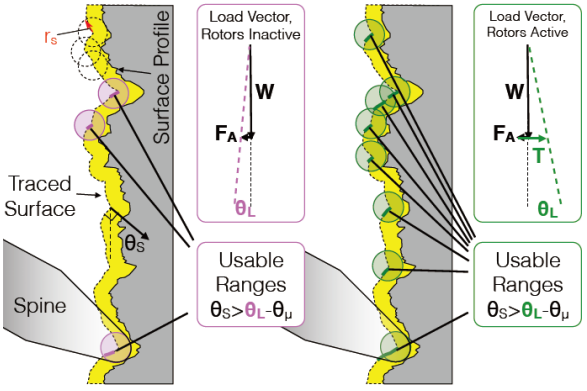


Fig. 7. With thrust pushing into the surface, the actual loading angle becomes shallower, and thus there are more usable asperities.

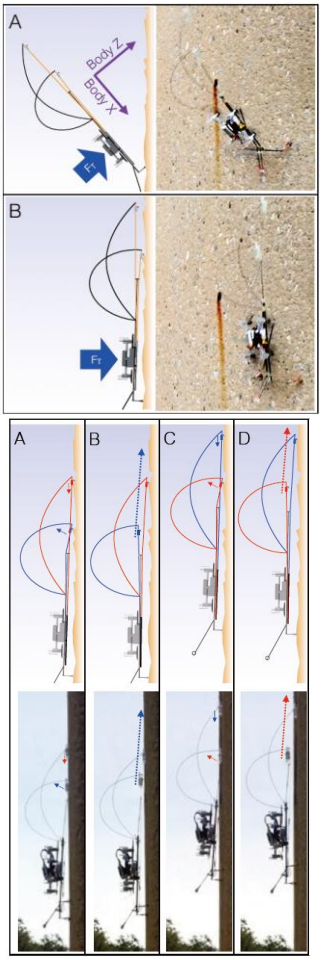


Fig. 8. SCAMP perching and climbing process.

be turned on yet because that force would keep pushing the quadrotor into the wall. As the quadrotor starts to fall, it also pivots backward on the tail spines because they are still engaged with the surface. The quadrotor then turns on full thrust to fly away. For inverted surface perching, the quadrotor flies directly into a surface to trigger the truss collapsing mechanism. It can detach from the surface by simply resetting the mechanism and falling. The average perching success rate on stucco and rough concrete surfaces is approximately 70%. The failed trials are primarily due to the local surface being too smooth or bad maneuvers due to wind disturbance.

Perching and climbing experiments on outdoor surfaces were conducted with the Stanford Climbing and Aerial Maneuvering Platform (SCAMP). The rotors can provide suction effects with forces approximately 1/4 of SCAMP's weight, while requiring only 10% of the maximum rotor power. As seen in Fig. 7, this is sufficient to triple the usable asperity density.

To perch, SCAMP flies tail-first into the wall. It detects the acceleration spike on impact and turns its rotors up to maximum. The tail acts as a pivot, and SCAMP attaches to the wall aerodynamically while its feet find good places to grip. If SCAMP misses a step, its onboard accelerometers detect the resulting free-fall and the rotors spool up to gently push the robot back onto the surface. Its feet then re-engage with the wall and it can resume climbing. When SCAMP is ready to resume flight, it rotates its servos to a special position (not seen during normal climbing) to deploy a "takeoff spine." When load is transferred from the engaged foot to this spine, it rotates SCAMP away from the surface. Once SCAMP has tipped back a few degrees, the rotors start to spool up to normal flight levels and the robot can resume aerial operation. The perching, climbing and take-off process is shown in Fig. 8.

A collaboration with UPenn aims at controlling a 500g quadrotor to perch on overhanging tubes and pipes (Fig. 9). The gripper has a bi-stable truss mechanism that can be reset by a servo. No sophisticated rebound spring is included in this design because the quadrotor flies directly to the pipe from underneath; thus there is no significant energy absorption needed. Preliminary results have shown that an increased wrist compliance is needed to account for misalignment between the gripper and the tubes. As the diameter of the tube becomes smaller, it is desirable to combine the gecko adhesives with solid fingers so that the fingers can wrap around the small tube and grasp more securely. An active rolling behavior of the quadrotor may be useful to preload the adhesives to accommodate misalignment.

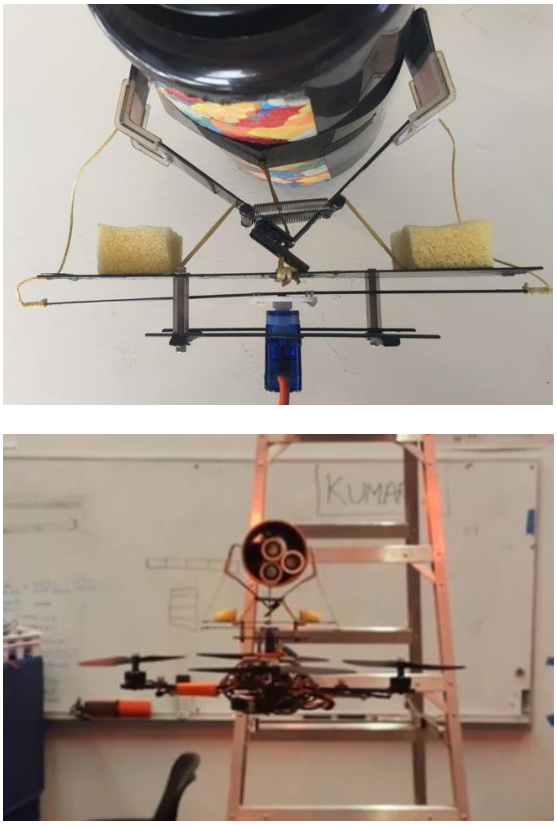
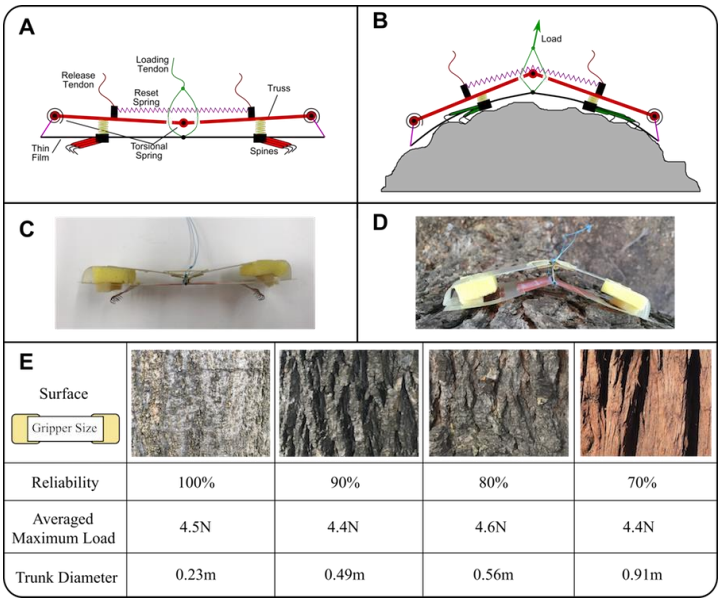


Fig. 9. **Top**: a bistable curved surface gripper attached to a water bottle. The servo can tension the release tendons to detach.



E Surface / Gripper Size				
Reliability	100%	90%	80%	70%
Averaged Maximum Load	4.5N	4.4N	4.6N	4.4N
Trunk Diameter	0.23m	0.49m	0.56m	0.91m

Fig. 10. Natural curved surface grippers made with micro-spines were evaluated on a variety of tree barks.

Flying and Tugging

All MAVs are limited by the maximum thrust they can produce. Utilizing a multimodal approach, we developed MAVs capable of both flight and forceful tugging employing controllable adhesives to generate contact forces with the environment. The resulting forces are over an order of magnitude

greater than the flier's bodyweight. A multimodal approach also addresses the limitation of separate classes of robots: strong microrobots have a hard time getting over obstacles [24]; flying robots are fast and maneuverable but weak. The multimodal result is a class of 100-gram robots mobile enough to navigate 3D environments, small enough to navigate within constrained spaces, and forceful enough to affect human-scale environments. An analysis of energetics, along with previous surveys of actuator scaling, suggest that slow forceful actuators benefit air vehicles increasingly as they become smaller, and can allow these robots to efficiently move a substantial variety of loads.

The combination of these capabilities suggests new possibilities for search-and-rescue robotics. An inspection application demonstrated in Fig. 11 shows a human operator hauling up a camera twice the weight of the 100 gram robot – a payload much too big for perching, let alone controlled flight.



Fig. 11. Flying, Micro, Tugging (FlyCroTug) robots were equipped with controllable adhesives and used to haul up payloads too heavy to perch with.

Conclusions and Future Work

In conclusion, we have developed new analytical and numerical models to describe micro-spine adhesion and gripper design. Perching systems have been equipped with both micro-spines and gecko adhesive grippers to perch on both flat and curved surfaces. Rotor assistive climbing has been shown to increase climbing robustness on outdoor rough surfaces and inertial control has been investigated to help perching systems better align with target surfaces and reduce rebound.

In the future, we will continue to push the limits of MAV perching and crawling in outdoor environments. The micro-spine curved surface gripper will be implemented on MAVs to perch on tree trunks and power poles and leave sensor packages on target surfaces. Distributed accelerometers and other onboard sensors will be investigated for improved perching failure detection and rapid recovery.

Acknowledgment

This work has been supported by MAST MCE-16-17.4.4 with additional support from NSF IIS-1161679 and NSF graduate fellowships for Estrada and Roderick.

References

[1] L. Daler, A. Klaptocz, A. Briod, M. Sitti, and D. Floreano, "A perching mechanism for flying robots using a fibre-based adhesive," in IEEE International Conference on Robotics and Automation, 2013.

[2] E. W. Hawkes, D. L. Christensen, E. V. Eason, M. A. Estrada, M. Heverly, E. Hilgemann, H. Jiang, M. T. Pope, A. Parness, and M. R. Cutkosky, "Dynamic surface grasping with directional adhesion," in IEEE International Conference on Intelligent Robot Systems, 2013.

[3] D. Mellinger, M. Shomin, and V. Kumar, "Control of quadrotors for robust perching and landing," in in Proc. Int. Powered Lift Conf, pp. 119–126, 2010.

[4] E. Glassman, A. Lussier-Desbiens, M. Tobenkin, M. Cutkosky, and R. Tedrake, "Region of attraction estimation for a perching aircraft: A lyapunov method exploiting barrier certificates," IEEE International Conference on Robotics and Automation, 2012.

[5] J.W.Roberts and R.Tedrake, "On the controllability of fixed-wing perching," American Control Conference, 2009.

[6] A.A.Paranjape, S.J.Chung, and J.Kim., "Novel dihedral-based control of flapping-wing aircraft with application to perching," IEEE Transactions on Robotics, 2013.

[7] H. Jiang, M. T. Pope, E. W. Hawkes, D. L. Christensen, M. A. Estrada, A. Parlier, R. Tran, and M. R. Cutkosky, "Modeling the dynamics of perching with opposed-grip mechanisms," in Robotics and Automation (ICRA), 2014 IEEE International Conference on, pp. 3102–3108, IEEE, 2014.

[8] J. Thomas, G. Loianno, M. Pope, E. W. Hawkes, M. A. Estrada, H. Jiang, M. R. Cutkosky, and V. Kumar, "Planning and control of aggressive maneuvers for perching on inclined and vertical surfaces," ASME 2015 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference, 2015.

[9] A. Kalantari, K. Mahajan, D. Ruffatto, and M. Spenko, "Autonomous perching and take-off on vertical walls for a quadrotor micro air vehicle," in Robotics and Automation (ICRA), 2015 IEEE International Conference on, pp. 4669–4674, IEEE, 2015.

[10] A. Lussier-Desbiens, A. T. Asbeck, and M. R. Cutkosky, "Landing, perching and taking off from vertical surfaces," The International Journal of Robotics Research, vol. 30, no. 3, pp. 355–370, 2011.

[11] M. Kovac, J. M. Germann, C. Hurzeler, R. Siegwart, and D. Floreano, "A perching mechanism for micro aerial vehicles," Journal of Micro-Nano Mechatronics, vol. 5, pp. 77–91, 2009.

[12] C. E. Doyle, J. J. Bird, T. A. Isom, J. C. Kallman, D. F. Bareiss, D. J. Dunlop, R. J. King, J. J. Abbott, and M. A. Minor, "An avian inspired

passive mechanism for quadrotor perching," Mechatronics, IEEE/ASME Transactions on, vol. 18, no. 2, pp. 506–517, 2013.

[13] Pope, Morgan T., *et al.* "A Multimodal Robot for Perching and Climbing on Vertical Outdoor Surfaces." *IEEE Transactions on Robotics* 33.1 (2017): 38-48.

[14] W. R. T. Roderick, Mark R. Cutkosky, and David Lentink. "Touchdown to take-off: at the interface of flight and surface locomotion." *Interface Focus* 7.1 (2017): 20160094.

[15] Libby, T., T. Moore, E. Chang-Siu, D. Li, D. Cohen, A. Jusufi and R. J. Full. 2012. Tail assisted pitch control in lizards, robots and dinosaurs. *Nature* 481, 181–184 (12 January 2012)

[16] Z. Dai, S. N. Gorb, and U. Schwarz, "Roughness-dependent friction force of the tarsal claw system in the beetle pachnoda marginata (coleoptera, scarabaeidae)," *Journal of Experimental Biology*, vol. 205, no. 16, pp. 2479–2488, 2002.

[17] A. T. Asbeck, S. Kim, M. R. Cutkosky, W. R. Provancher, and M. Lanzetta, "Scaling hard vertical surfaces with compliant microspine arrays," *The International Journal of Robotics Research*, vol. 25, no. 12, pp. 1165–1179, 2006.

[18] A. T. Asbeck and M. R. Cutkosky, "Designing compliant spine mechanisms for climbing," *Journal of Mechanisms and Robotics*, vol. 4, no. 3, p. 031007, 2012.

[19] H. Jiang, S. Wang, and M. R. Cutkosky, "Adhesion modeling of compliant microspines for rough surface grasping applications," *The International Journal of Robotics Research*, (submitted).

[20] E. W. Hawkes, et al. "Grasping without squeezing: Shear adhesion gripper with fibrillar thin film." *Robotics and Automation* (ICRA), 2015 IEEE International Conference on. 2015.

[21] M. A. Estrada, et al. "Free-flyer acquisition of spinning objects with gecko-inspired adhesives." *Robotics and Automation* (ICRA), 2016 IEEE International Conference on. IEEE, 2016.

[22] Roderick, W. R., Jiang, H., Wang, S., Lentink, D., & Cutkosky, M. R. (2017, July). Bioinspired Grippers for Natural Curved Surface Perching. In *Conference on Biomimetic and Biohybrid Systems* (pp. 604-610). Springer, Lecture Notes in AI.

[23] Lanzetta, M., and M. R. Cutkosky. "Shape deposition manufacturing of biologically inspired hierarchical microstructures." *CIRP Annals-Manufacturing Technology* 57.1 (2008): 231-234.

[24] Merz, Robert, et al. Shape deposition manufacturing. Engineering Design Research Center, Carnegie Mellon Univ., 1994.

[25] Christensen, David L., et al. "Let's all pull together: Principles for sharing large loads in microrobot teams." *IEEE Robotics and Automation Letters* 1.2 (2016): 1089-1096

Learning to Adapt Mobility to 3-D Unstructured Environments

MCE – 16-17 – 4.7

Pieter Abbeel, Gregory Kahn, Dave Held, Sergey Levine
Dept. of Electrical Engineering and Computer Sciences
UC Berkeley
Berkeley, CA, USA
pabbeel@berkeley.edu

***Abstract*—This project considers control for low-cost mobile (including aerial) platforms. Due to high variability in these platforms and the environments they might encounter, learning can play a critical role in establishing their capabilities. Traditionally reinforcement learning considers the problem of learning to control, but existing methods largely ignore safety during training. We propose two novel methods that directly account for safety during training (PLATO and ProbColl). We also propose self-supervised deep RL with generalized computation graphs for robot navigation.**

***Index Terms*—Reinforcement Learning, Safe Learning, Project Report for MAST - MCE 16-17 – 4.7**

I. Introduction

A. Background and Technical Challenges:

Low-cost platforms, such as the ones targeted in MAST research, tend to have higher variability in dynamics throughout their lifetime due to wear and tear and due to having to handle different environments. Moreover, each robot will have its own particular dynamics.

Reinforcement learning is a discipline concerned with learning to control a system through self-exploration (potentially starting from some prior knowledge). Reinforcement learning approaches have had varying levels of success, but the fundamental task of reinforcement learning --- autonomously learning to control a dynamical system --- has remained largely unsolved for robotic systems. One of the main challenges is that during execution of standard reinforcement learning algorithms the system might incur significant (maybe even irreparable) damage.

B. Research Objectives:

Making reinforcement learning more practical through consideration of safety at training time.

C. Technical Approach Summary:

We developed two novel approaches to reinforcement learning that explicitly account for safety during training:

- Policy Learning using Adaptive Trajectory Optimization (PLATO): *a method for training complex policies that combine perception and control by using a trajectory optimization teacher in the form of model-predictive control (MPC).* At training time, MPC chooses actions that make a tradeoff between succeeding at the task and matching the behavior of the current policy. By gradually adapting to the policy, MPC ensures that the states visited during training will allow the policy to learn good long-horizon performance. MPC makes use of full state information, which could be obtained, for example, by instrumenting the environment at training time. The final policy, however, is trained to mimic the MPC actions using only the observations available to the robot, which makes it possible to run the resulting policy at test time without any instrumentation. The algorithm requires access to at least a rough model of the system dynamics in order to run MPC during training, but does not require any knowledge of the observation model, making it feasible to use with complex, raw observation signals, such as images and depth scans. Since MPC is used to select all actions at training time, the algorithm never requires running a partially trained and potentially unsafe policy.
- Uncertainty-Aware Reinforcement Learning for Collision Avoidane (ProbColl): One of the central challenges in reinforcement learning is that a robot can only learn the outcome of an action by executing the action itself. Consider a robot learning to navigate a building while avoiding collisions. This scenario seemingly presents a quandary: the robot needs to learn how to avoid collisions in order to achieve the desired task, but to learn how to avoid collisions, the robot must experience (possibly catastrophic) collisions during training. The robot can overcome this quandary by first experiencing gentle collisions in order to learn about the environment; once the

robot is confident about the environment, the robot can avoid catastrophic failures in the future. Central to this approach is that the robot must be able to *reason about its own uncertainty* because highly-uncertain events are much more likely to cause catastrophic collisions.

- Self-supervised Deep Reinforcement Learning with Generalized Computation Graphs for Robot Navigation: Enabling robots to autonomously navigate complex environments is essential for real-world deployment. Prior methods approach this problem by having the robot maintain an internal map of the world, and then use a localization and planning method to navigate through the internal map. However, these approaches often include a variety of assumptions, are computationally intensive, and do not learn from failures. In contrast, learning-based methods improve as the robot acts in the environment, but are difficult to deploy in the real-world due to their high sample complexity. To address the need to learn complex policies with few samples, we propose a generalized computation graph that subsumes value-based model-free methods and model-based methods, with specific instantiations interpolating between model-free and model-based. We then instantiate this graph to form a navigation model that learns from raw images and is sample efficient. Our simulated car experiments explore the design decisions of our navigation model, and show our approach outperforms single-step and N-step double Q-learning. We also evaluate our approach on a real-world RC car and show it can learn to navigate through a complex indoor environment with a few hours of fully autonomous, self-supervised training. Videos of the experiments and code can be found at github.com/gkahn13/gcg

II. Related Work

A. Synergy with Existing Works and External Collaborations:

These ideas are applicable across a wide range of MAST platforms.

B. Technology Transitions Summary

N/A

C. Related Work in the Literature

In this work, we investigate how model-based reinforcement learning for robot collision avoidance can be made safe and reliable at both training and test time. Reinforcement learning has been applied to a wide range of robotic problems, ranging from locomotion and manipulation to autonomous helicopter flight [14], [4]. Model-free methods have been particularly popular due to their simplicity and favorable computational properties [24]. However, model- based methods are generally known to be more sample- efficient [3]. In this work, we adopt a model-based approach and learn an uncertainty-aware collision avoidance model; however, similar uncertainty estimation techniques could be extended also to model-free methods.

Several model-based robotic learning algorithms have been proposed that explicitly reason about uncertainty [3], [26]. Uncertainty estimates have been used to perform both risk-averse and risk-seeking, optimistic exploration [19]. The role of uncertainty estimation in our work is to avoid unsafe actions at training time until the model has gained sufficient confidence, which is largely orthogonal and complementary to prior work that seeks to improve exploration in order to accelerate learning. Combining these two directions is a promising direction for future work.

Uncertainty-aware model-based reinforcement learning has been explored in previous work using Bayesian models [25], [1]. While our work is similar in the overall aim, one of the central goals of our method is to directly process raw inputs from high-bandwidth sensors such as cameras, which necessitates the use of rich and expressive models, such as deep neural networks. Uncertainty estimation for deep neural networks is substantially more challenging, since these models are inherently discriminative. Recent work has proposed to use a Bayesian formulation of neural networks based on dropout [8], as well as to use the bootstrap for exploration [22] (but not, to the best of our knowledge, for uncertainty estimation for the purpose of safety). In this work, we demonstrate that combining both dropout and bootstrap can yield substantially improved uncertainty estimates for reinforcement learning tasks.

There is much prior work on safe robot control for safety-critical systems such as autonomous cars [29], legged robots [31], and quadrotors [20], [30], [9]. A number of recent works have sought to address the question of safety for learning-based robotic systems. Methods based on reachability provide appealing theoretical guarantees, but cannot cope with rich sensory input and are often difficult to scale to high-dimensional systems [23], [17], [10]. Several works have suggested using discriminative models, including neural networks, to learn safety predictors [2]. These methods generally take the approach of training a model to predict whether an unsafe action will occur, and reverting to a hand-designed safety controller if such a potential failure is detected. Our method offers two advantages over this approach. First, by directly estimating model uncertainty, we do not rely on a discriminative safety estimator. This approach is preferred in environments where the model might encounter previously unseen inputs because a discriminative safety estimator cannot provide meaningful predictions for completely novel inputs; in short, the discriminative safety estimator may erroneously conclude that an unsafe environment is safe. In contrast, a statistical uncertainty prediction such as bootstrapping is more likely to estimate high uncertainty in novel environments. Secondly, our approach does not assume the existence of a manually designed safety control, but instead naturally reverts to more cautious exploratory behavior in the presence of uncertainty. This makes the approach more automated, and does not require a safety mechanism that can recover from arbitrary unsafe situations.

We use deep neural networks to estimate the probability of collision from raw sensory inputs. Combining deep networks with reinforcement learning has been an active area of research in recent years, with applications to video game playing [18], control of simulated robots [27], [16], and manipulation [15]. However, most of these applications focus on task complexity or learning speed, rather than explicitly considering uncertainty and safety during training. Our previous work on autonomous flight [11] considers safety at training time by using model-predictive control (MPC) with ground truth state information. In contrast, our current work does not assume any access to ground truth state, which is advantageous for real-world deployment.

III. PLATO

A. Approach

We address the problem of learning control policies for dynamical systems, such as robots and autonomous vehicles. The system is defined by states x and actions u. The policy must control the system from observations o, which are in general insufficient for determining the full state x. The policy is a conditional distribution over actions πθ(u|ot), parametrized by θ. At test time, the agent chooses actions according to πθ(u|ot) at each time step t, and experiences a loss c(xt, ut). The next state is distributed according to the dynamics p(xt+1|xt, ut). The goal is to learn a policy πθ(u|ot) that minimizes the total expected cost.

In this work, we further assume that during training, our algorithm has access to the true underlying states x. This additional assumption allows us to use simple and efficient model-predictive control (MPC) methods to generate training actions. We do not require knowing the true states x at test time, since the learned policy πθ(u|ot) only requires observations. This training setup could be implemented in various ways in practice, including instrumenting the training environment (e.g. using motion capture to track a mobile robot) or using more effective hardware at training time (such as a more accurate GPS system), while only having access to cheaper and more practical hardware at test time. While this assumption does introduce some restrictions, we will show that it enables very efficient and relatively safe training, making it an appealing option for safety-critical systems.

We will train the policy πθ(u|ot) by mimicking a computational "teacher," rather than attempting to learn the policy directly with reinforcement learning. There are three key advantages to this approach: first, the teacher can exploit the true state x, while the final policy πθ is only trained on the observations o; second, we can choose a teacher that will remain safe and stable, avoiding dangerous actions during training; third, we can train the final policy πθ using standard, robust supervised learning algorithms, which will allow us to construct a simple and highly data-efficient algorithm that scales easily to complex, high-dimensional policy parametrization. Specifically, we will use MPC as the teacher. MPC uses the true state x and a model of the system dynamics (which we assume to be known in advance, but which in general could also be learned from experience). MPC plans locally optimal trajectories with respect to the dynamics, and by replanning every time step, is able to achieve considerable robustness to unexpected perturbations and model errors, making it an excellent choice for sample-efficient learning. Algorithm 1 outlines PLATO. We collect training trajectories by choosing actions ut according to an adaptive teacher policy π t λ (u|xt, θ), which is generated by optimizing the objective in Equation 1 at each time step via iLQG. We then update the learner policy πθ(u|ot) with supervised learning at the observations ot corresponding to the visited states xt to minimize the difference between πθ(u|ot) and the locally optimal policy, which is also obtained via MPC, but without considering the KL-divergence term. This approach ensures the teacher visits states that are similar to those that would be visited by the learner policy πθ, while still providing supervision from a near-optimal policy. Note that the MPC policy is conditioned on the state of the system xt, while the learned policy πθ(u|ot) is only conditioned on the observations. MPC requires access to at least a rough model of the system dynamics, as well as the system state, in order to robustly choose near-optimal actions. However, by training πθ on the corresponding observations, instead of the true states, πθ can learn to process raw sensory inputs without requiring true state observations, making it possible to run the learned policy with only the raw observations at test time. In the rest of this section, we describe the MPC teacher and the supervised learning procedure in detail.

Algorithm 1 PLATO algorithm

Initialize data $\mathcal{D} \leftarrow \emptyset$
for $i = 1$ **to** N **do**
  for $t = 1$ **to** T **do**
    Optimize π_λ^t with respect to Equation (1)
    Sample $\mathbf{u}_t \sim \pi_\lambda^t(\mathbf{u}|\mathbf{x}_t, \theta)$
    Optimize π^* with respect to Equation (2)
    Sample $\mathbf{u}_t^* \sim \pi^*(\mathbf{u}|\mathbf{x}_t)$
    Append $(\mathbf{o}_t, \mathbf{u}_t^*)$ to the dataset $\mathcal{D}$
    State evolves $\mathbf{x}_{t+1} \sim p(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t)$
  end for
  Train $\pi_{\theta_{i+1}}$ on $\mathcal{D}$
end for

B. Results

We evaluate PLATO on a series of simulated quadrotor navigation tasks. MPC is a standard choice for quadrotor control because approximate models are typically known in advance from standard rigid body physics and the vehicle specifications. However, effective use of MPC requires explicit state estimation and can be computationally intensive. It is therefore very appealing to be able to train an entirely feedforward, reactive policy to control a quadrotor performing navigation in obstacle-rich environments, directly in response to raw sensor inputs. During training, the vehicle might be

placed in a known, instrumented training environment to collect data using MPC, while at test time, the learned feedforward policy could control the aircraft directly from raw observations. This makes simulated quadrotor navigation an ideal domain in which to compare PLATO to prior work.

Performance of learned policies: In Figures 1a, 1b, 1d, and 1e, we present the mean time to failure (MTTF) of the learned policy $\pi\theta$ on the canyon and forest environments using the laser or camera sensors. The graphs show the MTTF of each policy at each iteration of the learning process, averaged over 10 training runs of each method with 20 repetitions each. Failure occurs when the quadrotor crashes into an obstacle, with the maximum flight time for each domain listed on the graphs. The results indicate that the PLATO algorithm is able to learn effective policies faster, and converges to a solution that is better than or comparable to the baseline methods. For some choices of β schedule and supervision scheme, some DAgger variants achieve similar final MTTFs, but always at a slower rate and, as discussed next, with significantly more training crashes.

Robustness during training: In Figures 1a, 1b, 1d, and 1e, we show the number of crashes experienced during training at each iteration. PLATO on average experiences less than one crash per iteration, comparable in performance to the baseline MPC method (supervised learning), indicating that mimicking the learner with a KL-divergence penalty does not substantially degrade the robustness of MPC. In contrast, both DAgger and coaching begin to experience a substantial number of failures when the mixing constant β drops. By carefully selecting the schedule for β, the number of crashes can be reduced. However, even with a carefully chosen schedule, the prior methods are vulnerable to non-stationary training environments, as illustrated in Figure 2c. In this experiment, the vehicle switches from the canyon to the forest halfway through training, and then switches back to the forest. PLATO still experienced on average less than one crash per episode in this mode, but the prior methods that directly execute $\pi\theta$ during training could not cope with this condition, since a policy trained only on the canyon cannot succeed on the forest without additional training. While this example might appear pathological, it is in fact a plausible training setup for a real quadrotor exploring a varied environment, such as different floors of a building. If the walls on one floor are painted, e.g., a different color than the rest, the learned policy could easily experience a catastrophic failure when entering the floor for the first time, even if it was consistently successful on preceding floors.

IV. Collision-Aware Exploration

A. Approach

The core component of our approach is an uncertainty-aware collision prediction model $P\theta$. Training this collision prediction model from experience presents a dilemma: the robot must first experience collisions in order to learn how to avoid collisions. We formulate a speed-dependent collision cost that uses uncertainty-aware collision estimates, resulting in the robot exploring cautiously when uncertainty is high and moving faster when uncertainty is low. This naturally arising behavior enables the robot to learn about collisions without experiencing catastrophic failures, and subsequently use these safe collision experiences to act more aggressively in the future.

An example application domain and desired application of the uncertainty-aware collision prediction model is the following: consider a quadrotor navigation task in which the objective is to fly fast and avoid collisions in an unknown environment. The quadrotor seeks to learn a collision prediction model that takes as input an image and a sequence of velocity commands and outputs the probability of collision. Initially, the quadrotor flies conservatively because the speed-dependent collision cost favors low-speed actions due to high uncertainty estimates of the collision prediction model. While flying conservatively, the quadrotor experiences safe collisions. These safe collisions, coupled with the associated images, are used to train the collision prediction model; the collision prediction model then learns how to associate images and velocity commands with the likelihood of colliding. As the algorithm continues and the collision prediction model uncertainty becomes low enough, the speed-dependent collision cost will favor high-speed flight.

Define abbreviations and acronyms the first time they are used in the text, even after they have been defined in the abstract. Abbreviations such as IEEE and SI do not have to be defined. Do not use abbreviations in the title or heads unless they are unavoidable.

Standard predictive neural network models are trained discriminatively, which means that, even though the network might achieve a high accuracy on samples drawn from the

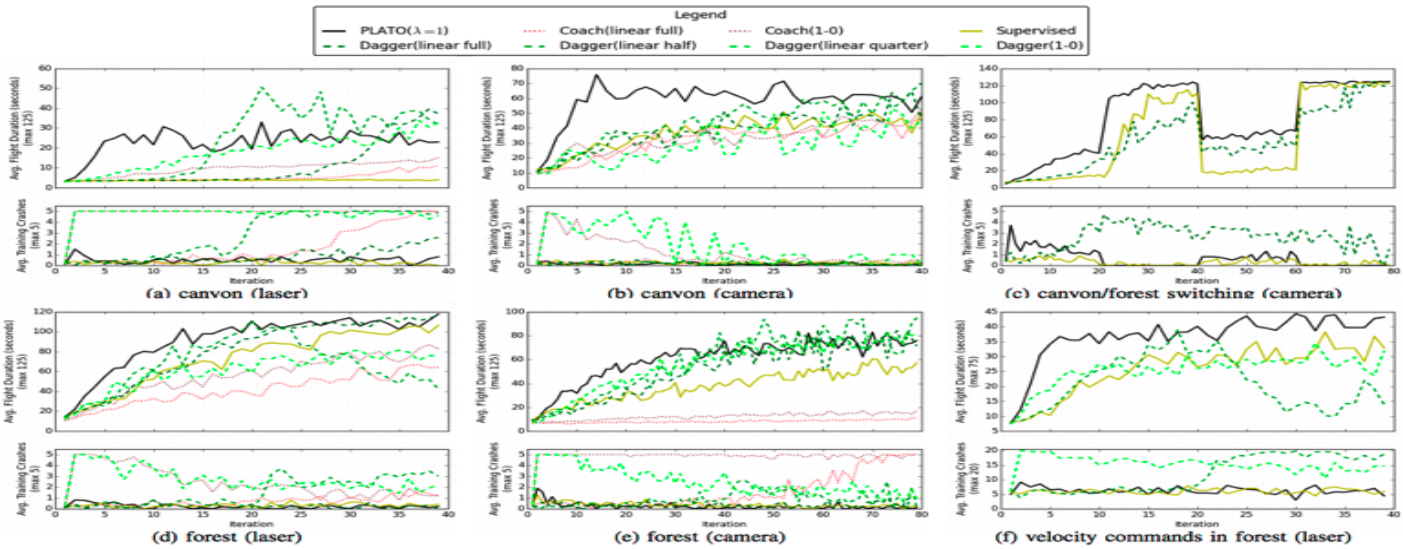


Figure 1: We compare PLATO to baseline methods in a winding canyon, a dense forest, and an alternating canyon/forest. For each scenario and learning methods, we trained 10 different policies using different random seeds. Each iteration required 2 minutes of flight time. Then for each policy, we evaluated the neural network policy trained at each iteration by flying through the scenario 20 times. Therefore each datapoint corresponds to 200 samples.

same distribution as the training data, it is very difficult to predict how the network would behave on data drawn from a different distribution. While it is possible to train a neural network model that outputs a mean and a variance as its prediction, this model is not in general guaranteed to output high variances for unfamiliar inputs because the network is by definition trained only on the datapoints that are in the training set. Indeed, such a method for estimating variance is only effective at estimating the inherent noise in the data, and the variance estimates are not a meaningful indication of the model's own uncertainty about its predictions. To produce accurate uncertainty estimates for data that is outside of the training distribution, we must explore techniques that go beyond direct discriminative training. In order to obtain accurate uncertainty estimates from our model, we use two techniques: bootstrapping and dropout.

B. Results

We evaluate our uncertainty-aware collision prediction model in a simulated environment consisting of a cylindrical obstacle of radius 0.2m (Fig. 5). The quadrotor's objective CTASK is to fly forward at 0.5m/s. Unless stated otherwise, the time horizon H = 6. Each discrete time step corresponds to $\delta = 0.2$ seconds, therefore the planning horizon is δH seconds.

We evaluate different parameter settings of our method and a baseline along two dimensions. The first dimension is the success rate, which tells us for what fraction of flights the vehicle was able to successfully navigate around the obstacle.

To evaluate the benefits of uncertainty-aware exploration, we can also examine the collisions that the vehicle experiences during training. Since low-speed collisions are typically less harmful, we evaluate the different methods in terms of a second dimension: the total kinetic energy absorbed by the vehicle during training, which is proportional to the sum of the squared velocities during all collisions.

Fig. 2 compares different parameter settings for the uncertainty-aware neural network collision prediction model in terms of their effect on the total kinetic energy and the success rate. In each comparison, we vary one parameter of the model. The standard parameter setting, marked as a star, corresponds to λSTD = 1, 5 bootstraps, and a dropout percentage of 0.05. We compare against a baseline model that is not risk-averse, which corresponds to parameters λSTD = 0, 5 bootstraps, and a dropout percentage of 0.05; the baseline model still uses bootstrapping and dropout in order to be a fair baseline because bootstrapping and dropout often improves the performance of neural networks (which is their typical use case). Note that the model using the standard parameter settings experiences fewer high-speed collisions compared to the baseline while still maintaining comparable task performance. This indicates that uncertainty- aware collision prediction is effective in reducing the number of training-time collisions.

We provide a preliminary evaluation our approach in a real-world environment consisting of a single obstacle, in which the objective is to fly around the obstacle while avoiding collisions (Fig. 1). We ran our experiments using a Parrot Bebop 2 quadrotor. We used the ROS bebop_autonomy package, which allows the laptop to send linear velocity commands and receive the onboard images in real-time.

The goal of this experiment was to evaluate whether our proposed uncertainty-aware model could make risk-averse predictions from real-world images. Due to time constraints, we did not run the full model-based RL algorithm, but instead trained our model on a dataset consisting of 112 trajectories, each with 8 datapoints, collected via teleoperation. In our experiments, we varied the parameter λSTD, which controls the degree to which the model seeks to minimize risk. The results, presented in Fig. 4, compare the quadrotor behavior with λSTD = 0 (Fig. 4a) and with λSTD = 1 (Fig. 4b). The results indicate that larger values of λSTD result in low-speed collisions while smaller values of λSTD result in high-speed collisions. Our ongoing experiments aim to evaluate the full

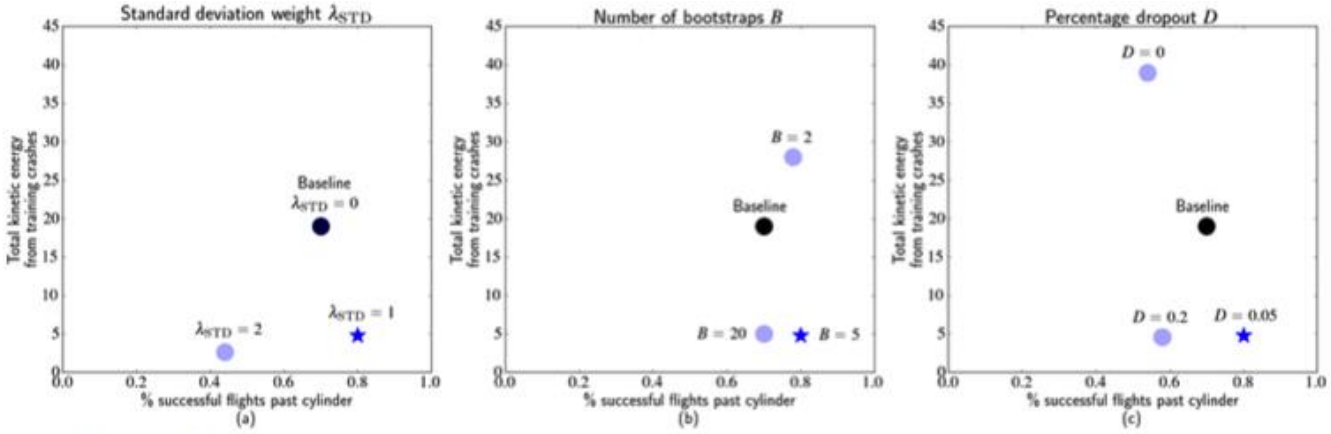


Fig. 2: **Collision prediction model comparisons**: We investigated the effect of changing three parameters (a) the standard deviation weight λ_{STD} from Eqn. 1 (b) the number of bootstraps B and (c) the percentage dropout D with respect to the total kinetic energy experienced from crashes during training and the percentage of flights that successfully flew around the cylindrical obstacle. The best performance method has parameters $\lambda_{\text{STD}} = 1$, $B = 5$, and $D = 0.05$ and is marked with a star. The baseline method does not account for uncertainty and therefore has parameters $\lambda_{\text{STD}} = 0$, $B = 5$, and $D = 0.05$. For all other methods, $\lambda_{\text{STD}} = 1$, $B = 5$, and $D = 0.05$ *unless a parameter is explicitly set according to the plot.* In all plots, points closer to the bottom right are better.

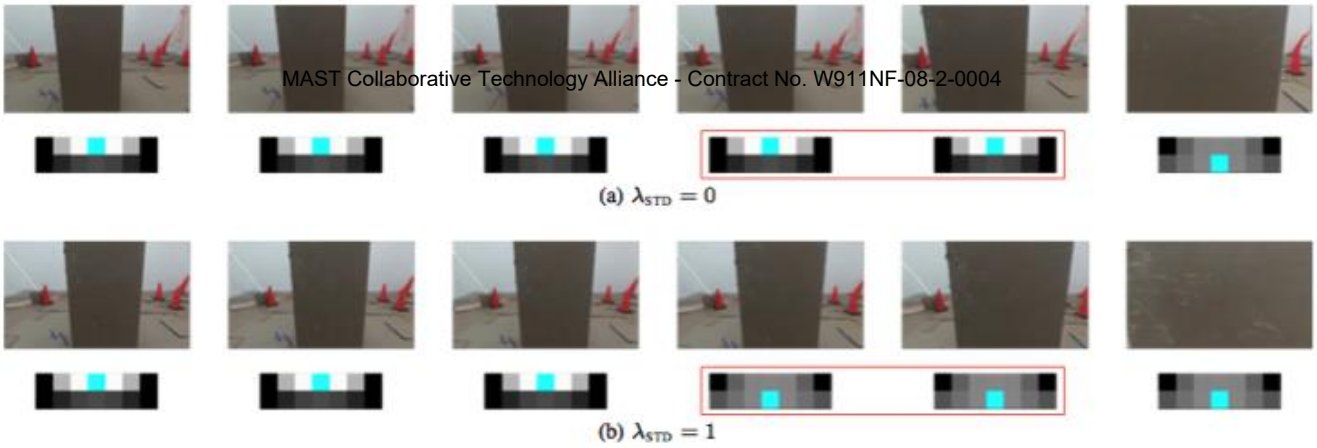


Fig. 4: **Real-world quadrotor experiments**: We compare the effect of varying λ_{STD}. Given an incompletely trained collision prediction model, we run MPC with (a) $\lambda_{STD} = 0$ and (b) $\lambda_{STD} = 1$. For each setting of λ_{STD}, we show images captured by the quadrotor (top row) and a heatmap indicating the cost of different actions MPC is choosing between (bottom row). Each cell in the heatmap corresponds to a heading direction (indicated by the x-axis) and a speed (indicated by the y-axis); the color corresponds to the cost associated with executing that action, where black is high cost, white is low cost, and cyan corresponds to the least cost action. The heatmaps indicate the quadrotor crashes into the obstacle at high-speed in (a) while crashing into the obstacle at low-speed in (b). Although the model does not correctly predict which action will avoid the obstacle, the risk-averse estimate correctly deduces that the model is in a region of high uncertainty (outlined in red), and our collision-aware cost automatically causes the quadrotor to slow down.

RL algorithm on the physical hardware, and to evaluate a larger variety of parameter settings. Videos of the real- world flights can be found in the supplementary materials and online: https://sites.google.com/site/probcoll/

V. Generalized Computation Graphs for Robot Navigation

A. Approach

Learning offers considerable promise for mobile robotic systems: by observing the outcomes of navigational decisions in the real world, mobile robots can continuously improve their proficiency and adapt to the statistics of natural environments. Not only can learning-based systems lift some of the assumptions of geometric reconstruction methods, but they offer two major advantages that are not present in analytic approaches: (1) learning-based methods adapt to the statistics of the environments in which they are trained and (2) learning-based systems can learn from their mistakes. The first advantage means that a learning-based navigation system may be able to act more intelligently even under partial observation by exploiting its knowledge of statistical patterns. The second advantage means that, when a learning-based system does make a mistake that results in a failure, the resulting data can be used to improve the system to prevent such a failure from occurring in the future. This second advantage, which is the principal focus of this work, is closely associated with reinforcement learning: algorithms that learn from trial-and-error experience.

Reinforcement learning methods are typically classified as either model-free or model-based. Value-based model-free approaches learn a function that takes as input a state and action, and outputs the value (i.e., the expected sum of future rewards). Policy extraction is then performed by selecting the action that maximizes the value function. Model-based approaches learn a predictive function that takes as input a state and a sequence of actions, and output future states. Policy extraction is then performed by selecting the action sequence that maximizes the future rewards using the predicted future states. In general, model-free algorithms can learn complex tasks, but are usually sample-inefficient, while model-based algorithms are typically sample-efficient, but have difficulty scaling to complex, high-dimensional tasks.

We explore the intersection between value-based model-free algorithms and model-based algorithms in the context of learning robot navigation policies. Our work has three primary contributions. The first contribution is a generalized computation graph for reinforcement learning that subsumes value-based model-free methods and model-based methods. The second contribution is instantiations of the generalized computation graph for the task of robot navigation, resulting in a suite of hybrid model-free, model-based algorithms. The third contribution is an extensive empirical evaluation in which we: (a) investigate and discover design decisions regarding our robot navigation computation graph that are important for stable and high performing policies; (b) show our approach learns better policies than state-of-the-art reinforcement learning methods in a simulated environment; and (c) show our approach can learn a successful navigation policy on an RC car in a real-world, indoor environment from scratch using raw monocular images.

B. Results

We demonstrated the efficacy of our approach on a 1/10th scale RC car in a real-world environment. The task was to navigate the 5th floor of Cory Hall at 1.2m/s, which is challenging as these hallways contain tight turns, changing illumination, and glass walls.



Figure 5: Real-world RC car experiments on the 5th floor of Cory hall. The path the robot can follow is drawn on the floor plan (left), however, this path is not provided to the RC car. Three example trajectories of the RC car navigating with our learned policy are

The RC car learning system was set up to maximize time spent gathering data, minimize computational burden for the car, and be fully autonomous. The computer onboard the RC car is an NVIDIA Jetson TX1, which is intended for embedded deep learning applications. However, all model training is performed offboard on a laptop for computational efficiency, while model inference is performed onboard. We therefore made the system fully asynchronous: the car continuously runs the reinforcement learning algorithm and sends data to the laptop, while the laptop continuously trains the model and periodically sends updated model parameters to the car. For full autonomy, the car automatically detects collisions using the onboard IMU and wheel encoder, and automatically backs up if a collision occurs. The only human intervention required is if the car flips over, which occurred approximately every 30 minutes.

In evaluating our approach, we chose the best design decisions from our simulation experiments: the model outputs are collision probabilities trained using classification, a large model horizon (H=12, corresponding to 3.6m lookahead), and no bootstrapping. All other settings were the exact same as the simulation experiments.

From training with only 4 hours of data in a complex, real-world environment using only raw camera images and no prior knowledge, the car could navigate significant portions of the environment. For example, the best trajectory travelled 197m, corresponding to nearly 2 loops (and 8 hallway lengths). However, sometimes the policy fails; additional training should correct these mistakes.

We compared our method with a random policy and double Q-learning trained using the data gathered during our approach's learning. We trained double Q-learning in this way in order to compare the performance of the algorithms given the same state distribution. Our approach travels 17 times farther than the random policy and 7 times farther than double Q-learning. Qualitatively, our approach was smoothly driving straight when possible, while double Q-learning was exhibiting bang-bang control in that its steering angle is always at the limits.

VI. Conclusions and Further Work

We presented:
(i) PLATO, an algorithm for learning complex, high-dimensional policies that combine perception and control into a single expressive function approximator, such as a deep neural network. PLATO uses a trajectory optimization teacher to provide supervision to a standard supervised learning algorithm, allowing for a simple and data-efficient learning method. The teacher adapts to the behavior of the neural network policy to ensure that the distribution over states and observations is sufficiently close to the learned policy, allowing for a bound on the long-term performance of the learned policy. Our empirical evaluation, conducted on a set of challenging simulated quadrotor domains, demonstrates that PLATO outperforms a number of previous methods, both in terms of the robustness and effectiveness of the fina; and (ii) a model-based method for learning obstacle avoidance strategies that uses uncertainty estimates to auto- matically generate safe exploration strategies. Our method is based on predicting the probability of collision conditioned on raw sensory inputs and a sequence of actions, using deep neural networks. This predictor can be used within a model-predictive control pipeline to choose actions that avoid collisions with high probability. In regions of high uncertainty, our risk-averse cost function naturally causes the robot to revert to a cautious low-speed exploration strategy, without any explicit manual engineering of safety controllers or fail-safe mechanisms. We evaluate our method in simulation on a range of synthetic scenarios designed to evaluate the safe exploration strategy, compared to a method without uncertainty estimates, and test a variety of uncertainty estimation techniques. We also demonstrate that our method can produce effective safe exploration behavior on a simulated obstacle avoidance task, and provide preliminary results showing that the model can be used to control the collision speed in a real-world quadrotor obstacle avoidance task.

(iii) Self-supervised Deep Reinforcement Learning with Generalized Computation Graphs for Robot Navigation: Enabling robots to autonomously navigate complex environments is essential for real-world deployment. We propose a generalized computation graph that subsumes value-based model-free methods and model-based methods, with specific instantiations interpolating between model-free and model-based. We then instantiate this graph to form a navigation model that learns from raw images and is sample efficient. Our simulated car experiments explore the design decisions of our navigation model, and show our approach outperforms single-step and N-step double Q-learning. We also evaluate our approach on a real-world RC car and show it can learn to navigate through a complex indoor environment with a few hours of fully autonomous, self-supervised training.

A promising direction of future work is to combine our method with optimistic exploration strategies that mitigate this sort of "uncertainty- aware panic," for example by choosing to gradually back off on the uncertainty penalty when the vehicle repeatedly fails to make progress.

Another promising direction for future work is to generalize our approach beyond collision prediction to other model-based reinforcement learning scenarios. The principle of uncertainty-aware prediction of future events can be readily applied to any feature of the environment, including the expected cost, and exploring this extension to general reinforcement learn- ing problems could produce effective and safe exploration techniques for a wide range of robotic scenarios.

More detailed write-ups on both contributions are available here:

(i) https://arxiv.org/abs/1603.00622
(ii) https://arxiv.org/abs/1702.01182
(iii) https://arxiv.org/abs/1709.10489

REFERENCES

[1] F. Berkenkamp, A. Krause, and A. Schoellig, "Bayesian optimization with safety constraints: Safe and automatic parameter tuning in robotics," in arXiv:1602.04450, 2016.
[2] S. Daftry, S. Zeng, J. A. Bagnell, and M. Hebert, "Introspec- tive perception: Learning to predict failures in vision systems," in arXiv:1607.08665, 2016.
[3] M. Deisenroth and C. Rasmussen, "Pilco: A model-based and data- efficient approach to policy search," in ICML, 2011.
[4] M. P. Deisenroth, G. Neumann, and J. Peters, "A survey on policy search for robotics," in Foundations and Trends in Robotics, 2013.
[5] B. Efron and R. Tibshirani, "The jackknife, the bootstrap and other resampling plans," in SIAM, 1982.
[6] ——, An introduction to the bootstrap. CRC press, 1994.
[7] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in arXiv:1506.02142, 2015.
[8] Y. Gal, R. Mcallister, and C. Rasmussen, "Improving pilco with bayesian neural network dynamics models," in Data-Efficient Machine Learning workshop, ICML, 2016.
[9] J. Gillula and C. Tomlin, "Guaranteed safe online learning via reachability: Tracking a ground target using a quadrotor," in ICRA,2012.
[10] ——, "Reducing conservativeness in safety guarantees by learning disturbances online: Iterated guaranteed safe online learning," in RSS, 2012.
[11] G. Kahn, C. Zhang, S. Levine, and P. Abbeel, "Plato: Policy learning using adaptive trajectory optimization," in arXiv:1603.00622, 2016.
[12] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in ICLR, 2015.
[13] A. Kleiner, A. Talwalkar, P. Sarkar, and M. I. Jordan, "The big data bootstrap," in ICML, 2012.
[14] J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey," vol. 32, no. 11, 2013, pp. 1238–1274.
[15] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," 2016.
[16] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," in arXiv:1411.0247, 2015.
[17] A. Majumdar and R. Tedrake, "Funnel libraries for real-time robust feedback motion planning," in arXiv:1601.04037, 2016.
[18] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and R. M., "Playing atari with deep reinforcement learning," in Workshop on Deep Learning, NIPS, 2013.
[19] T. Moldovan, S. Levine, M. Jordan, and P. Abbeel, "Optimism-driven exploration for nonlinear systems," in ICRA, 2015.
[20] M. Mueller and R. D'Andrea, "Stability and control of a quadrocopter despite the complete loss of one, two, or three propellers," in ICRA, 2014.
[21] V. Nair and G. Hinton, "Rectified linear units improve restricted boltzmann machines," in ICML, 2010.
[22] I. Osband, C. Blundell, A. Pritzel, and B. Van Roy, "Deep exploration via bootstrapped dqn," in NIPS, 2016.
[23] T. Perkins and A. Barto, "Lyapunov design for safe reinforcement learning," in JMLR, 2002.
[24] J. Peters and S. Schaal, "Policy gradient methods for robotics," in IROS, 2006.
[25] C. Richter, W. Vega-Brown, and N. Roy, "Bayesian learning for safe high-speed navigation in unknown environments," in ISRR, 2015.
[26] J. Schneider, "Exploiting model uncertainty estimates for safe dynamic control learning," in NIPS, 1997.
[27] J. Schulman, S. Levine, P. Moritz, M. I. Jordan, and P. Abbeel, "Trust region policy optimization," in ICML, 2015.
[28]Srivastava,G.Hinton,A.Krizhevsky,I.Sutskever,andR.Salakhutid-nov, "Dropout: A simple way to prevent neural networks from overfitting," in JMLR, 2014.
[29] C. Urmson and et. al., "Autonomous driving in urban environments: Boss and the urban challenge," in Journal of Field Robotics, 2008.
[30] M. Watterson and V. Kumar, "Safe receding horizon control for aggressive MAV flight with limited range sensing," in IROS, 2015.
[31] P. Wieber, "Viability and predictive control for safe locomotion," in IROS, 2008.

Robophysical approach to discovery of principles of effective interaction of ambulatory robots with multicomponent environments

Project: MCE-16-2.1

Jeffrey J. Aguilar, Henry C. Astley, Yasemin O. Aydin, Yang Ding, Nick Gravish, Christian M. Hubicki, Andras Karsai, Chen Li, Shengkai Li, Ryan Maladen, Hamidreza Marvi, Nicole Mazouchova, Ben McInroe, Daria Monaenkova, Feifei Qian, Jennifer M. Rieser, William Savoie, Perrin E. Schiebel, Sarah Sharpe, Xiaobin Xiong, Tingnan Zhang, Paul Umbanhowar, Daniel I. Goldman

School of Physics
Georgia Institute of Technology
Atlanta, Georgia USA
daniel.goldman@physics.gatech.edu

***Abstract*—Ambulatory robots will be required to traverse unpredictable multi-component heterogeneous substrates whose constituent parts vary in stiffness, flowability, orientation and gap sizes. As a specific example of a challenging complex terrain, consider substrates containing granular media (GM) with particle sizes spanning multiple orders of magnitude, with different compactions, and irregularly distributed obstacles buried within. While enhanced robot locomotion on or within homogeneous flowable substrates has been realized by our group through an interplay of laboratory techniques and modeling, during our MAST funding, comparable techniques did not yet exist for locomotion on heterogeneous, flowable, terrestrial substrates. To test robotic locomotion and navigation and extract informative models on heterogeneous substrates, a new set of experimental apparatuses were built to parameterize and systematically prepare complex terrain, as well as new robotic locomotion platforms (i.e. snakes, hexapods, monopods, bipeds, salamanders and mudskippers). The result was a demonstrated ability to describe and predict the features of successful locomotion for these MAST-type robots on irregular terrain. Terrain with exaggerated heterogeneities, such as boulders and solid pegs, revealed exotic scattering dynamics similar to wave interference; deformable sheets yielded governing equations typically associated with general relativity. This report presents this progression of capabilities and concepts, and illustrates how it flowed from MAST-funded work.**

***Index Terms*—Robophysics, granular media, locomotion, substrate physics, heterogeneous and complex terrain**

I. Introduction

A major challenge in creating the next generation of mobile robots is expanding the scope of predictive models for vehicle-terrain interaction (e.g., terramechanics [6, 36]), from large tracked and treaded vehicles on homogeneous ground to arbitrarily-shaped and actuated locomotors moving on and within complex heterogeneous terrestrial substrates. This challenge has spurred the creation of new computational tools, new and unique robotic platforms, and a deeper theoretical understanding of successful movement through complex terrain.

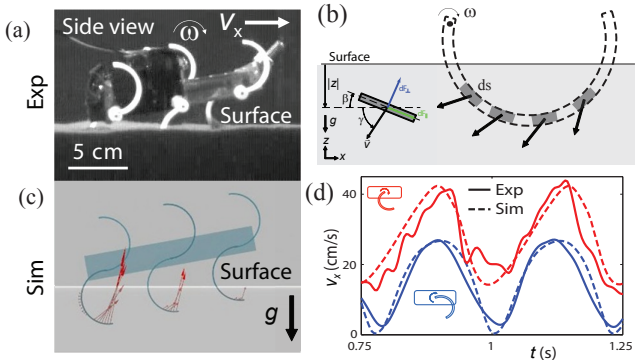


Fig. 1: Resistive Force Theory applied to legged locomotion on dry granular media. (a) RHex-like robot (Xplorer) robot during alternating tripod gait locomotion. (b) Illustration of the basic idea behind Resistive Force Theory (RFT) for movement of a RHex c-leg into granular media. Each infinitesimal element on the intruding leg is characterized by its tangent direction (or normal direction) and its velocity v, which are used to calculate the ground resistance. In true fluids, these forces can be described by Stokes' law, while for granular media, they are measured in experiment. (c) Simulation of Xplorer locomotion using RFT. Red arrows indicate granular reaction forces at each segment. (d) Comparison of forward speed vs. time between an experimental robot and corresponding RFT simulation using c-legs and reversed c-legs. Adapted from [19].

Initially, this research program developed an understanding and computational framework for enabling a better understanding multi-legged locomotion on complex homogenous substrates, such as granular media [12]. Resistive force theory models enabled locomotion of monopedal, bipedal, hexapedal, salamander- and mudskipper-like robots on flowable terrain. This work evolved to envelop motion synthesis tools for

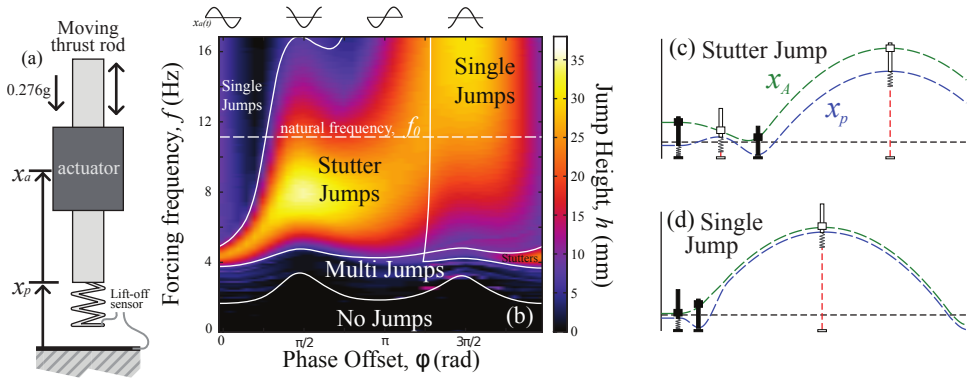


Fig. 2: Robotic hopping strategies on hard ground. (a) Diagram of SLIP-inspired jumping robot diagram (b) Color map of experimental jump heights vs forcing frequency and phase offset during 1 cycle sine-wave forcing. Approximately 20,000 experiments are represented. White lines separate regions of different jump types. Panels (c) and (d) illustrate trajectories for the stutter and single jumps [2].

robotic control and expanded into the increasingly challenging regime of heterogeneous terrain. Our results show how collisions with heterogeneous obstacles, such as single boulders, can produce scattering effects with attractive and repulsive features that are useful for anticipatory control [27, 28]. Further, multiple peg obstacles yield wave-interference-like scattering effects on the robots, e.g. robot snakes. This emergent diffraction pattern is often observed with quantum objects, like photons and electrons, but manifests here on robotic scales. Similarly, we exploited these collision-effects to create an amoeba-like robot, which locomotes via the asymmetry in self-deforming active internal particles, or *smarticles*. This trajectory of advances demonstrates the feasibility of understanding and exploiting the structure and dynamics of terrain to enable a diversity of locomotion modes in complex environments, as well as focus on fundamental principles.

II. Prior Work and Work during Funding

A. RFT and Terradynamics

The theories of aero- and hydrodynamics predict animal movement and guide device design in air and water through the computation of lift, drag, and thrust forces. Although models of terrestrial legged locomotion have focused on interactions with solid ground, many animals move on substrates that flow in response to intrusion. However, locomotor-ground interaction models on such flowable ground are often unavailable. Our work was the first to develop the principles of locomotion on homogeneous granular media, a terradynamics of flowable substrates [19].

We developed a force model for arbitrarily-shaped legs and bodies moving freely in granular media [37], and used this granular resistive force theory (RFT) to predict a small legged robot's locomotion on granular media using various leg shapes and stride frequencies, as illustrated in Fig. 1. Our study revealed a complex but generic dependence of stresses in granular media on intruder depth, orientation, and movement direction and gives insight into the effects of leg morphology and kinematics on movement.

B. Hard Ground Hopping

While jumping is commonly performed by a variety of complex robotic and biological locomotors and is commonly viewed as a component of animal running gaits, simple spring-mass jumpers can exhibit rich dynamics and provide insight into more complex systems. To gain an understanding of the dynamics of various jumping strategies, we performed systematic experiments on a simplified 1-D SLIP hopper illustrated in Fig. 2a. Balance issues were avoided by constraining the robot with a vertical air bearing. The actuation strategy of the motor was systematically varied by using a simple sinusoidal control template and adjusting parameters such as phasing and forcing frequency. The robot's simplicity facilitated the complete automation of the experiment, allowing kinematic data to be collected for 20,160 jumps over a periods of days.

Two stereotyped jumping strategies emerged: the single jump, characterized by a single push-off, and the stutter-jump consisting of a preliminary hop followed by a push-off (Fig. 2c,d). A simple analytical model of the system revealed how optimal timing parameters of such maneuvers were influenced by resonance, transient dynamics and hybrid dynamic transitions between aerial and ground phases (particularly in the case of stutter jumps) [2]. Examining the simplest form, the SLIP model, revealed rich dynamics that are intrinsic to jumping performance, lessons that would have proved challenging to learn with more complicated robots.

C. Granular Hopping

The locomotion of terrestrial animals or robots is typically studied in scenarios where unyielding environments deform the locomotor or where non-compliant locomotors deform yielding substrates. However, in many robotically and biologically relevant situations, such as impulsive interactions during running and hopping, actively forced compliant appendages interact with complex deformable substrates like granular media. To investigate such interactions, we studied a jumping robot that performed a sequence of fast impulsive and forced intrusions in granular media with the objective of high speed take-off [1] (Fig. 3a). Non-forced impact studies suggest that

such one dimensional intrusion forces can be predicted with a relation that accounts for both quasi-static frictional forces as well as velocity dependent inertial drag [17]. To test these force relations, we tested the robot (originally used in hard ground jumping experiments [2]) in a bed of granular media that controls volume faction (loose packed to close packed) via air fluidization, air pulses, and bed shaking. A separate motor raised the jumper between experiments so that the granular media's state could be reset. After hundreds of automated experimental jumps, comparison between jump heights in experiment and simulation revealed that models combining inertial drag and hydrostatic friction were insufficient, inspiring a more in-depth examination of granular kinematics [1].

This required a simultaneous study of both robot and substrate, which was achieved by moving the jumper next to the sidewall for direct measurement of grain kinematics via PIV analysis of high speed video capture, as Fig. 3c depicts. This analysis revealed a cone of grains that solidified underneath the foot and moved with the foot. A geometric model of this cone's growth with respect to the foot's intrusion depth in conjunction with empirical RFT measurements [19] was able to predict the quasistatic resistive force on the foot (Fig. 3d). Furthermore, the cone's dynamics signaled that an added mass effect - an effect observed in fluids in which the mass of an intruder moving through a media is effectively increased as it accelerates and displaces surrounding material - was producing a non-negligible impact on the robot's kinematics. Consequently, a 1-D reactive force theory incorporating added mass was able to accurately reproduce experimental jump heights in simulation [1].

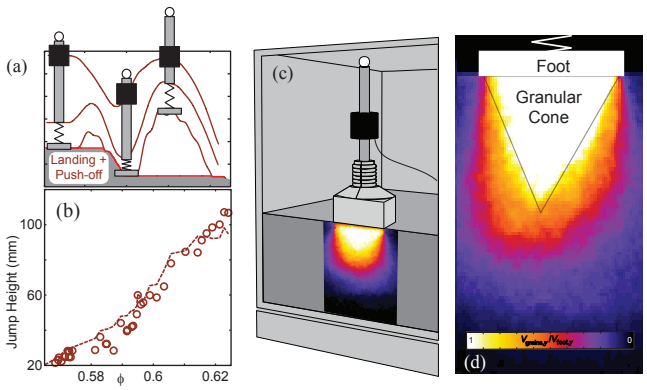


Fig. 3: Jumping on granular media. (a) Diagram of a stutter jump in time. Motor pulls rod up for an initial hop and immediately pushes off yielding granular media upon the first landing. (b) A simulation (dashed) of the stutter jump incorporating added mass into the reaction force model of the granular media was performed using experimental forcing trajectories at different granular volume fractions and exhibited good agreement with experiment (circles). (c) This granular added mass model was intuited and confirmed from analysis of poppy seed grain kinematics during jumping, where the robot was moved up to the clear sidewall and highspeed video captured grain flow during jumping. (d) PIV velocity field of sidewall grain flow reveals a cone of material that jams under the foot and effectively makes the foot more massive, producing an added mass effect. Figure adapted from [1].

D. *FlipperBot*

Flipper-based terrestrial locomotion (like that used by sea turtles) highlights the challenge of using aquatically adapted appendages to locomote on land. Understanding this mode of locomotion has benefited from a robophysics approach, which illuminates the influence of mechanical design on the induction of qualitatively different substrate behaviors, in this case generating solid states of granular media to improve locomotor performance and reduce suspectibilty to failure. Our biological field observations of Loggerhead hatchlings [23] revealed that the hatchlings utilize granular solidification by bending their wrist during limb-ground interaction with sand, and employ a rigid wrist while walking on hard ground, allowing for comparable forward speeds on both surfaces. To understand how such limb-ground interactions affect locomotion performance we created FlipperBot a servo-motor-driven robot which propels itself using limbs and flat-plate flippers [24].

With a fixed-wrist joint, FlipperBot was prone to failure by interacting with ground disturbed by previous steps, resulting in a decrease of flipper generated forces as the per-step distance was reduced. Conversely, a flexible wrist made the robot less prone to failure since there was less interaction with previously disturbed ground. The use of FlipperBot additionally allowed for systematic changes in flipper insertion depth, showing how extremely shallow penetration depths (1.4 mm) prevented forward movement. Overall, FlipperBot proved essential in understanding how kinematic control and structural design affect the dynamics of flipper-based locomotion on sand. This flexible limb use enables control of the state of the granular media, is remarkably similar to how sidewinders modulate body contact length to control the granular state.

E. *Muddybot*

In the evolutionary transition from aquatic to terrestrial environments, early tetrapods faced the challenges of terrestrial locomotion on flowable substrates, such as sand and mud of variable stiffness and incline. The morphology and range of motion of appendages is revealed in fossils; however, biological and robophysical studies of modern taxa have shown that movement on such substrates can be sensitive to small changes in appendage use. Using a biological model (the mudskipper) [26], a physical robot model (the MuddyBot) (Fig. 5E) inspired by skeletal reconstructions of early terrestrial walkers such as the Acanthostega (Fig. 5B), granular drag measurements, and theoretical tools from geometric mechanics, we demonstrated how tail use can improve robustness to variable limb use and substrate conditions. We hypothesize that properly coordinated tail movements could have provided a substantial benefit for the earliest vertebrates to move on land. In identifying these effects, we are elucidating minimal templates for terrestrial locomotion on complex terrain.

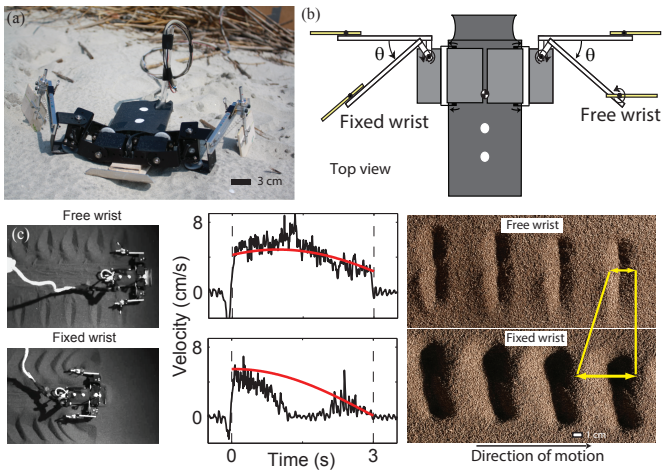


Fig. 4: Flipper-based terrestrial locomotion. (a) Sea-turtle inspired FlipperBot. (b) Diagram indicating fixed (left) and flexible (right) wrist flipper configurations. The robot inserts flippers into media to some depth and then strokes from front to back. (c) Comparison of forward displacement of FlipperBot with a flexible wrist (top) and rigid wrist (bottom). Experimental (black) and theoretical (red) velocity vs. time in center panel. Fixed wrist induces a granular frictional fluid which produces a larger area of granular disturbance than free wrist (right panel), thus increasing the amount of weakened substrate that subsequent strokes interact with. The free wrist model is kinematic since the flipper-generated forces remain below the granular yield stress. Adapted from [24].

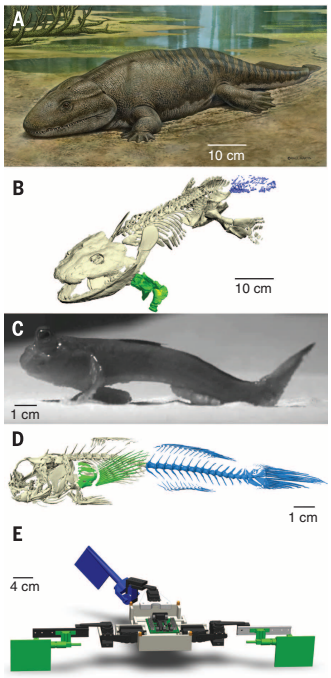


Fig. 5: Minimal template for terrestrial locomotion on flowable surfaces. (A) A reconstruction of Ichthyostega (∼ 360 million years ago), an example of an early tetrapod body plan, by Raul Martin. (B) Skeletal reconstruction of Ichthyostega, an example of an early tetrapod body plan [from (20)], highlighting the pectoral limbs (green) and tail (blue). (C) The mudskipper (Periophthalmus barbarus), a biological model for early terrestrial loco- motors. (D) A microcomputed tomography scan reconstruction of a mudskipper skeleton, highlight- ing the pectoral fin (green) and tail (blue). (E) The MuddyBot, a 3D printed robot developed to model the locomotion of crutching early tetrapods. Limbs are in green and the tail is in blue.

F. Sidewinders

Many limbless locomotors inhabit sandy deserts, and the biological and neuromechanical principles by which they move effectively in dry granular materials are largely unknown. To better understand strategies for successful locomotion on granular materials, we studied the above-surface locomotion of the Sidewinder rattlesnake, *C. cerastes*, a snake that uses a peculiar form of locomotion called sidewinding (Fig. 6A). In this gait, the animal translates through successive lifting of body segments while other segments remain in static (rolling/peeling) contact with the ground. We find that snake employs an anti-slip strategy to negotiate granular inclines, and can ascend surfaces even close to the angle of repose. Implementing this control strategy in a limbless robot (in collaboration with Howie Choset's group at Carnegie Mellon University, see Fig. 6B) enables the robot to ascend sandy inclines with comparable ability to the animal [22]. More broadly, these robotic tests reveal that the snakes movement can be described by a mixture of waves of propagating body curvature in two orthogonal planes. We can model the maneuverability of the snake (and improve the robot's locomotion efficacy) by spatial and phase modulations of this two-wave template [5], see Fig. 6C.

G. SCATTER

For future robots to effectively maneuver across terrains of interest such as disaster sites, comets and extraterrestrial planets, we must understand the mechanisms of effective movement in such terrains. This will only be possible through the parametrization and systematic creation of heterogeneous granular ground. However, characterizing how the nearly limitless variability in size, shape, granular compaction, gap, orientation, etc. affects locomotion can seem like a daunting task to probe by hand. In fact, without automation, collecting the large statistical data sets that will likely comprise the type of analysis that robophysics provides will be nearly impossible. The following study exemplifies automation's utility in studying locomotion on heterogeneous substrates.

We [28] took the first steps in this ambitious direction with the development of a fully-automated terrain creation system called Systematic Creation of Arbitrary Terrain and Testing of Exploratory Robots. The SCATTER system consists of an air fluidized bed that sets sand compaction, two tilting actuators that control the substrate inclination, and a universal jamming gripper [7] that retrieves and re-distributes boulders (objects of similar size to the robot) and the robot after each locomotion test. Using the SCATTER system, properties of heterogeneous substrates such as compaction, orientation, obstacle shape/-

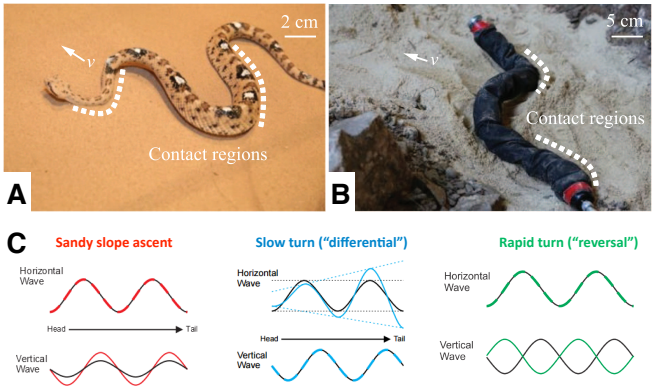


Fig. 6: Template modulation for maneuvering while sidewinding. (A) A Sidewinder rattlesnake (*C. cerastes*) locomoting on a granular substrate of 20 inclination. (B) Carnegie Mellon University Modular snake robot (courtesy of Howie Choset) traversing granular media at 20 inclination. The arrow indicates the direction of motion. (C) Examples of motion that can be produced with slight variations in amplitude modulation and phasing of the vertical and horizontal waves.

size/distribution, and obstacle mobility within the substrate, can be precisely controlled and varied. A high speed tracking system was also integrated into the SCATTER system, to track the robot kinematics in 3D, observe robot and boulder interaction, and locate robot and boulders after each test for gripper retrieval. SCATTER, which is high throughput like the above mentioned systems ($\sim$ 200 tests/day without human intervention) has already revealed important properties of open loop locomotor scattering of a small hexapedal robot (15 cm, 150 g) during various boulder interactions. The SCATTER trackway was filled with fine granular media (the "sand") and a single "boulder" (3D printed convex objects of different geometries) was located on the surface. Analysis of the robot's trajectory indicated that each interaction could be modeled as a scattering event with attractive and repulsive features, whose magnitude depended sensitively on the local boulder surface inclination at the initial contact point. Depending on the contact position on the boulder, the robot will be scattered in different directions. For a larger heterogeneous field with multiple boulders, the trajectory of the robot can be statistically estimated using a superposition of the scattering angle from each boulder. This scattering superposition can be applied to a variety of heterogeneity, including different geometry, orientation, and roughness.

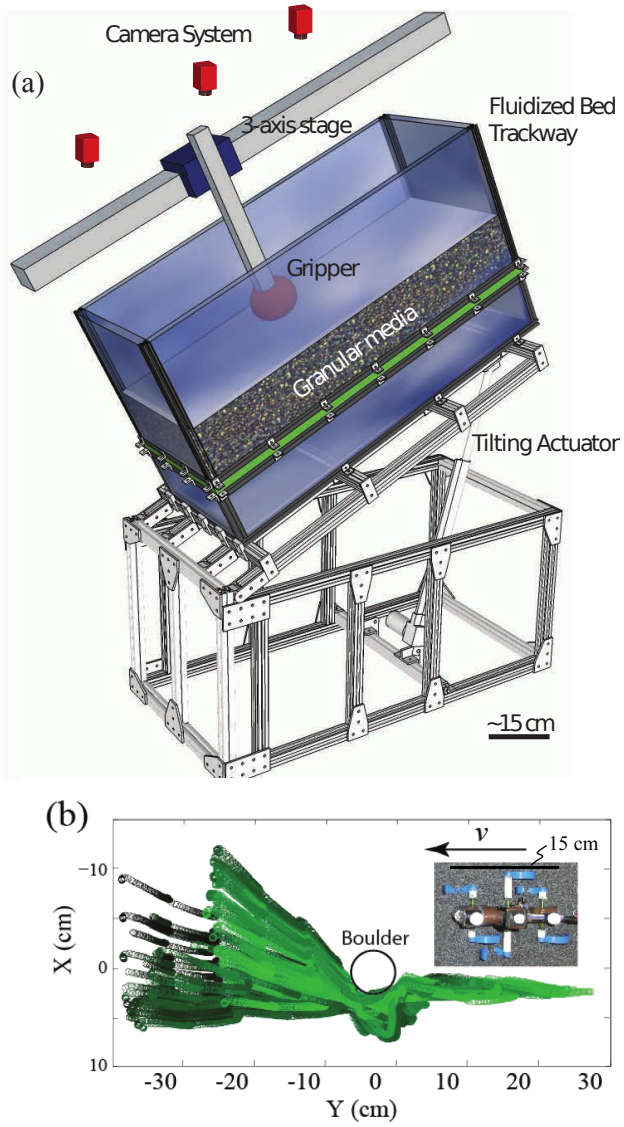
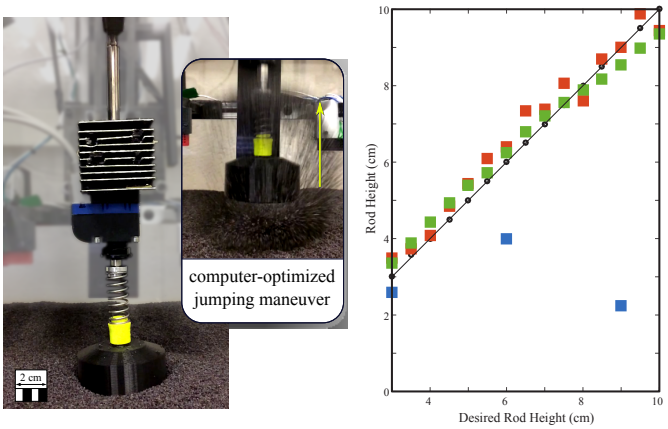


Fig. 7: Locomotion with single boulder collisions. (a) The Systematic Creation of Arbitrary Terrain and Testing of Exploratory Robots (SCATTER) system allows for automated three axis manipulation of heterogeneities such as boulders as well as robots traversing a granular substrate with specified ground stiffness and surface incline. (b) Robot interaction with single spherical boulder. The center of the boulder is set as the origin. Robot (inset) center of mass trajectories for the single boulder interaction. Trajectory colors represent variation in robot initial fore-aft positions, which influenced the angle of scatter caused by the boulder. Various initial positions were automated with the SCATTER system. Figure adapted from [27].

III. Current Work

A. Optimal Jumping

In pursuit of dynamic walking and running robots, significant progress has been made toward control and motion planning with the nonlinear, hybrid dynamics of legged machines. However, it is less common to study a robot's interactions with the dynamics of the terrain beneath its feet. The world is full of dirt, sand, mud, and other deformable terrain, but by and large, underlying physical models employed in motion planning typically assume rigid contact. Here, we present a tractable method for motion planning subject to rapid interactions with a complex substrate, specifically loose-packed poppy seeds (a granular medium which has proven useful in the study of complex-terrain locomotion). We experimentally validate our planning on a testbed spring-legged jumping robot (Fig. 8).



(a) (b)

Fig. 8: Optimal jumping on a yielding substrate. (a) A robotic jumping platform performing an optimized jump on loosely packed poppy seeds, where (b) we can predictively generate optimal jumping maneuvers which accurately achieve desired hop heights, despite the relative complexity of granular media [15] (Green: Hard ground strategy (HG) on HG, Blue: HG strategy on granular media (GM), Red: GM strategy on GM).

To set an aggressive benchmark to test the effective accuracy of our motion plans, we explored the problem of open-loop impulsive jumping on loose-packed grains.

This work presented [15] and validated a tractable and accurate approach to motion planning through granular media (e.g. sand and dirt) using an added-mass model [1]. Specifically, we tested impulsive jumping on loosely packed poppy seeds to see if we could quickly (1 sec) compute open-loop jumping trajectories which achieve desired apex heights. Our ability to generate optimal jumps on poppy seeds using the added-mass model (3.9 ± 2.1 mm, 6% error) was nearly as accurate in hitting target heights as our jumps on hard ground (3.4 ± 1.5 mm). We believe this is a success of the model and motion planning formulation, especially considering the complex media and open-loop control. Further, using hard ground strategies on granular media was not at all successful, demonstrating the utility of thoughtfully considering the terrain dynamics when generating robot control. Ongoing work integrates the use of machine-learning methods to learn terrain dynamics on-the-fly [8].

B. Salamander Robot

To understand how vertebrates made the transition from water to land, we first analyzed the motion of mudskippers, modern analogs of early tetrapods, and showed that their tail is helpful for land locomotion, especially on sandy slopes [25]. Now we have begun to work on salamanders, another existing animal model for tetrapod evolution. We aim to show theoretically and experimentally that coordinated body bending may have improved the robustness of locomotion on challenging substrates especially on inclines. To investigate mechanisms governing effective locomotion of salamander, we have built a robophysical model and started to test its performance on hard ground and granular media. This open-loop servo-driven 3D-printed robot (450 g, 38 cm long) has four limbs, an actuated trunk, and an active tail (Fig. 9A).

Each limb has two servo motors to control the vertical position and the step size of the limb. A joint in the middle of the body controls horizontal bending. The design of and materials that make up the robot allow us to easily change morphological parameters of the body. Using a fully-automated setup [29], we have performed experiments to systematically test the importance of body morphology, limb coordination, tail use, body flexibility, and control in locomotion on yielding substrates, and to compare successes and failures of the robot with the animals.

Preliminary robot experiments varied the body angle amplitude, α, from 0 to 90° while keeping the limb angle relative to the body, β, fixed and assessed robot performance by measuring the displacement per step of the robot in granular material (poppy seeds, grain size 1 mm). Fig. **??** shows step length averaged over 10 cycles for a symmetric diagonal gait, a reasonable approximation for salamanders on hard ground. We see that step length monotonically increases as body-bending amplitude increases, with a maximum step length at $\alpha = 90°$. We have also used geometric mechanics [18], a tool for mapping local body deformations of an object (joint angles) to the lab frame taking into account the material response properties of the environment (in collaboration with Professor Howie Choset at Carnegie Mellon University). With this model, a largest displacement for a corresponding body angle can be identified by the integral of the height function, which is highlighted as the sinusoidal curve in Fig. 9B. Predicted displacements exhibit a quantitatively similar trend to the experiments for all body-bending amplitudes, Fig. 9C.

C. Snake Scattering

Limbless locomotors such as snakes gracefully navigate through a wide range of varied terrain, from tall grasses to debris-littered forest floors to the yielding sandy substrates of the desert. Despite the seeming simplicity of their movement, the interaction with the ground coupled with intermittent obstacle collisions can give rise to complex dynamics. To begin to understand the mechanics and control underlying even relatively simple locomotor-obstacle interactions, we built an open-loop position-controlled 13-segment snake-like robot from 12 servo motors connected together with 3D-printed plastic brackets. Joint angles are controlled by prescribing servo motor positions to follow a “serpenoid” curve [21]. Wheels affixed to the bottom of each segment introduce a frictional anisotropy with the ground and create a strongly preferred direction of motion, enabling the robot to convert the local joint angle displacements into translational motion. We study the dynamics of this robot moving through a model heterogeneous terrain in which a row of five evenly-spaced pegs, oriented perpendicular to the robot’s initial direction of motion, extend vertically upward from a firm, flat surface. A top-down view of the setup is shown in Fig. 10A. To sample all possible interactions, the robot begins each experiment in the same configuration and is placed at different positions within

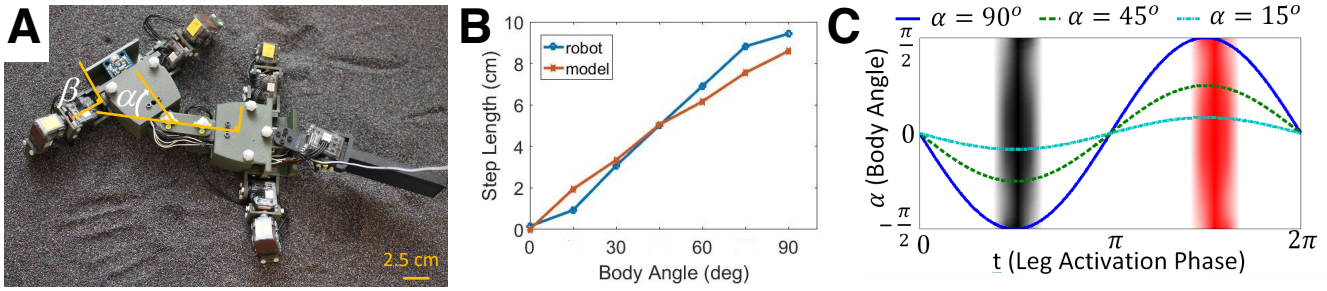


Fig. 9: Geometric mechanics and legged locomotion. (A) Robophysical model of a salamander. (B) Step length vs. body angle amplitude for experiment (blue) and geometric mechanics prediction (red). (C) The height function corresponding to a sinusoidal gait (dark blue $\alpha = 90°$, green $\alpha = 45°$ and light blue $\alpha = 15°$).

a region with lateral and longitudinal dimensions set by the peg spacing and distance traveled in one undulation cycle, shown as the green box in Fig. 10A. After this initialization, positions of markers atop each segment are recorded as a function of time as the robot travels toward, interacts with, and subsequently exits the peg array. Two example trajectories in Fig. 10A show the sensitivity of the final trajectory on the initial placement of the robot.

We find that, despite the uniformly-sampled initial conditions, the trajectories along which the robot emerges from the peg array are not uniform and that the robot only travels along a few preferred paths, which we characterize by the angle of rotation of the direction of travel, as indicated by θ in Fig. 10A. Points within the green box indicate experimental initial conditions for 329 experiments, with color indicating the final heading of the robot. To visualize all possible interactions, we look at the summation of trajectories originating within the green box in Fig. 10A, normalized by the total number of experiments. This "heatmap" (Fig. 10B) shows the fraction of experiments for which that pixel in the image is visited by the robot, with brighter colors indicating higher probability. Heatmaps for four spacings are shown, and in all cases, the locations visited by the robot after the interaction are not uniform, with the robot preferentially following a small number of trajectories. The robot deflections become smaller as peg spacing increases, see Fig. 10B. Numerical simulations are in agreement with experiments and allow for a more thorough exploration of this dependence, which we find is surprisingly well-fit by the functional form expected for wave fields and subatomic particles.

Finally, we have begun to investigate how to alter the setup to allow the robot to continue to travel along its original path. Instead of infinitely-massive pegs that cannot displace, simulations enable us to easily vary the mass of the pegs, which allows them to move (in the transverse direction only) in response to a collision with the robot, see schematic in Fig. 10C. Heatmaps of simulated robot trajectories for a few different peg masses, M, are shown in Fig. 10D. As the mass of the pegs is decreased, the existence of these preferred paths is diminished and eventually destroyed as the pegs are allowed to move more in response to the collision with the robot. Future explorations will consider allowing the robot to locally deform its body in response to a collision, which we hypothesize may alter the final trajectory statistics and make different locations in the environment more accessible.

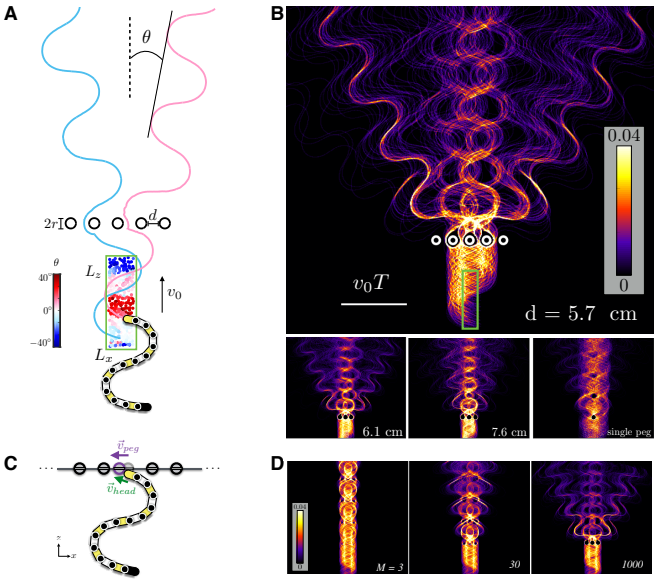


Fig. 10: Mechanical diffraction in limbless robots. (A) Experimental setup for robotic snake scattering experiments. A snake robot travels toward, interacts with, and exits the peg array. Two trajectories of the robot head are shown. Colors within the green box indicate the scattering angle, θ, for each initial placement. (B) Heatmaps of robot head-trajectories for four spacings. The robot deflections become smaller as peg spacing increases. (C) Simulation setup for finite-mass pegs that are now allowed to move (in response to a collision with the robot) in the direction transverse to the initial heading of the snake. (D) The resulting heatmaps for moving pegs as a function of the mass of the pegs for d = 6.1 cm. Large mass approaches the rigidly-affixed case, which can be seen by comparing the right heatmap to the bottom left heatmap in (B). As mass decreases, the pegs are easier to push out of the way and the snake deflects less.

D. Robotic Ant Foraging

This work investigated foraging strategies for excavators in confined pathways, e.g., foraging ants. A small swarm of

decentralized, autonomous robotic ant analogues were constructed which operate in a narrow tunnel to repeatedly transport granular media from an excavation site to a deposit site. High-population robotic swarms have the advantage of being robust to individual robotic failures, while also increasing productivity. However, there is a reduction in productivity when the population grows beyond the capacity of the work space, at which point the robots become preoccupied with obstacle-avoidance maneuvers rather than productive tasks. Interference is a common issue in swarm robotics which arises when the operating demands of a population of robots, such as power or work space, exceed the available resources [9, 38]. This spatial interference is a factor that biological organisms such as fire ants have developed methods to avoid; by applying behavioral protocols based on observations of these organisms we are able to enhance the overall performance of the robotic system.

The robotic test bed is able to capture complex agent interactions that take place as the population of robots increases. Given a large population with an active protocol in which each robot acts industriously to maximize individual productivity, the entire swarm will coalesce into an irrecoverable traffic jam, thereby hindering overall productivity. Behaviors which have been observed in biological ants and tested in simulation, such as the Lorenzian and reversal strategies, have been implemented and effectively modulate the competition for limited space in the tunnels. Populations which would readily congest given an active protocol will instead exhibit a steady flux of robots within the tunnel and the productivity of the swarm will be retained in spite of the overpopulated state of the system. By incorporating behaviors which reduce competition for space, the collective swarm behavior remains optimal even at high robot populations while also maintaining high robustness to robot malfunctions.

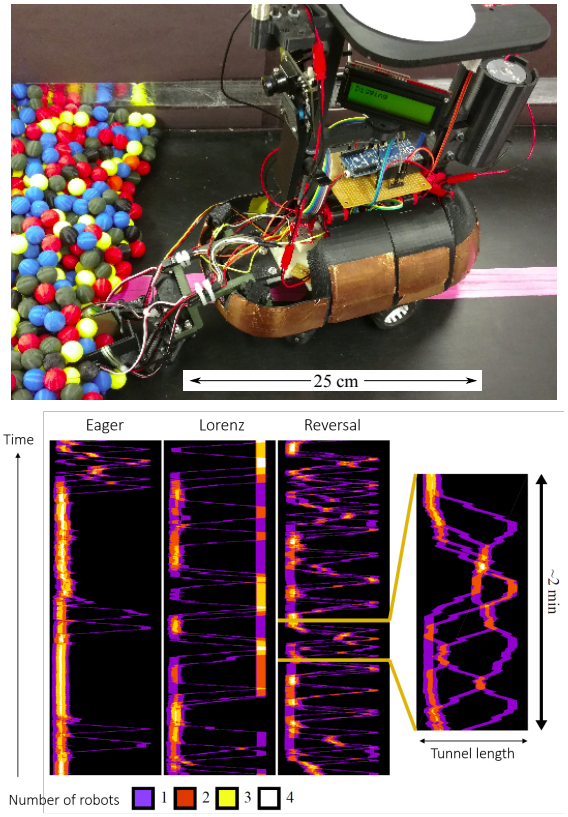


Fig. 11: Collective excavation in confined conditions. (top) A robophysical platform for ant foraging experiments, excavating magnetic spheres. (bottom) Space-time plots showing number of ant robots (indicated by the colors) within the tunnel for different excavation strategies. The excavation site is on the left and deposit site is on the right. These plots highlight the increase in robot flow throughout the tunnel that arises when all robots are not equally eager to dig.

E. Deformation-Induced Locomotion

Previous studies of locomotion have focused on the interactions between locomotors and dissipative substrates. We explore the dynamics of locomotion resulting from movement on a deformable, elastic membrane. In this environment, the mass of a single robot is enough to locally deform the surface, providing a way for multiple robots to affect each other through the deformations they induce on surface on which they move. To characterize the way in which the motion of the car is affected by the deformation of the membrane, we have designed an experiment in which a spandex sheet is stretched over a round, 2.5-m-diameter frame. A linear actuator is attached to the fabric at the center of the ring, allowing the creation of a central depression in the membrane, see Fig. 12A. A photograph of our robot is self-propelled, differential-drive car on the elastic membrane is shown in Fig. 12B. The motion of the car is recorded as a function of time and as a function of the depth of the central depression, and we find that the car tends to be attracted to a closed orbit when the depth is large (upper left trajectory in Fig. 12C). For small depths, the car precesses along elliptic orbits (lower right trajectory in Fig. 12C).

Remarkably, the precession angles, $\Delta\phi$, of the car moving on the membrane are well described by the Schwarzschild metric solution, from which it is expected that $\Delta\phi \sim \frac{1}{A(1-e^2)}$, where A is the half major-axis and e is the eccentricity. To calculate A and e from the precessed orbit, we use the distance of the aphelion and perihelion from the center r_{min} and r_{max} and we find that $A = \frac{1}{2}(r_{max} + r_{min})$, $e = \frac{r_{max}-r_{min}}{r_{max}+r_{min}}$. For a closed orbit, the winding angle of the trajectory between two consecutive r_{max} or r_{min} should be 2π. For a precessing orbit, the difference from 2π of the angle is the precession angle $\Delta\phi$. Measured precession angles are compared to Schwarzschild-prediction solutions in Fig. 12C.

F. Bipedal Locomotion

Legged robots, especially bipedal humanoids, enjoy significant advantages over machines with wheels in locomoting both through unpaved natural terrains and in constructed human environment. Little is known about robotic bipedal walking on soft granular terrains such as sand, gravel, dirt, etc. Well-developed hard ground walking schemes such as Zero Moment Point (ZMP) [16] and Hybrid Zero Dynamics (HZD) [4, 35]

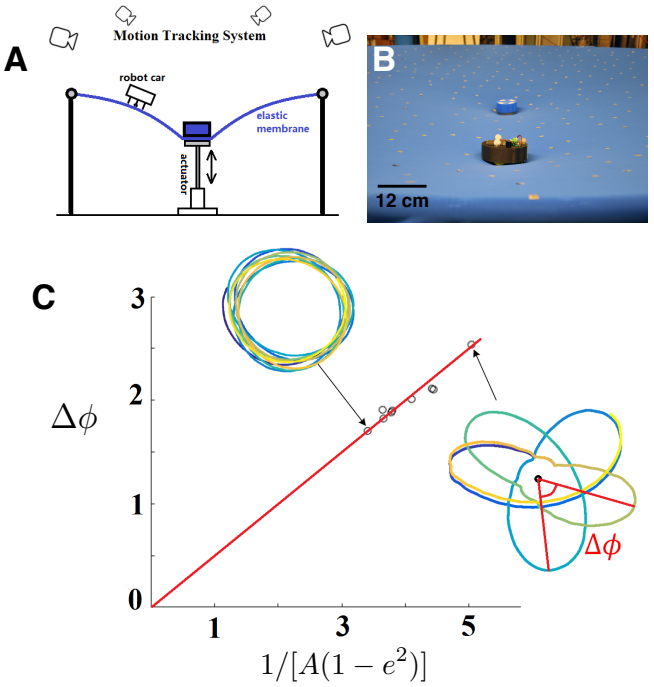


Fig. 12: Locomotion on an elastic membrane. (A) Experimental setup. (A) Differential-drive car resting on the membrane. (C). Predicted precession vs measured precession angle, along with two example trajectories: an orbiting car (upper left) and a precessing car (lower right). The functional form is what is expected for the Schwarzschild solution in General Relativity.

have shown success in numerous robotic systems, but fail once the rigid contact assumption is removed. To overcome the challenges of bipedal walking on deformable terrains, we analyze the foot-terrain interactions of flat-footed robotic walking on homogeneous loosely packed granular terrain (in collaboration with Aaron Ames).

The robot platform is shown in Fig. 13A-B. Using a newly discovered granular force overshoot phenomenon during paused intrusion and granular resistive force theory (RFT) [19], we model a static stability region for the center of mass (CoM) projection of an inverted pendulum on the granular surface. The stability region, shown in Fig. 13C, is experimentally validated by using a simple Foot-Table apparatus 13C. More importantly, we demonstrate that the stability region approximates the admissible zero moment point (ZMP) region for walking with a constant CoM height. We experimentally validate this admissible ZMP region using preview control [16]. Additional foot placement strategy for traditional ZMP walking is necessary to accommodate the deformation of the terrain. By interpreting the stability region as the maximum reaction moment of the terrain at different depth, we can also formulate the hybrid dynamics of walking on granular terrain and generate a walking motion via trajectory optimization techniques. We experimentally validate the walking gaits generated from Partial Hybrid Zero Dynamics (PHZD), showing that the stability region criterion yields stable dynamic walking on deformable granular terrain.

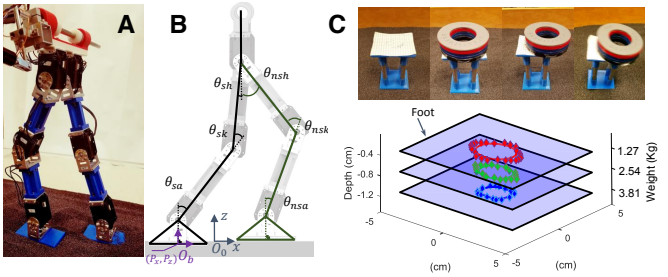


Fig. 13: Bidepal walking on granular material. (A) A photo of the bipedal robot. (B) The planar biped and coordinates. (C) (top) Mapping the static stability region of a foot in poppy seeds is determined by placing a mass at different locations around the foot. When the center-of-mass is too far away from the center-of-geometry of the foot, the foot becomes unstable and falls over. (bottom) The static stability region at different intrusion depth for the robot's foot dimensions.

G. Wheeled Locomotion

Our tools allow us to investigate the mechanics of single wheeled locomotion in loosely packed granular media. A recent study [34] has demonstrated that granular Resistive Force Theory (RFT) [19] successfully models the movement of slowly rotating wheels in the quasistatic limit, predicting that slip, S, is 1.2 and insensitive to rotational speed, ω. To test if RFT applies at increased ω, we study a 16 cm diameter wheel in a granular bed of loosely packed poppy seeds (1 mm diameter) driven at constant ω from 15 to 360 deg/s. As in previous work, for low ω, S is constant. However, as ω increases beyond 90 deg/s, S increases as the wheel loses traction with the granular media. At the highest ω, the wheel fails to move due to a rapid excavation of material underneath. Our results suggest that granular RFT may have to be modified to capture the new effects observed at higher ω.

H. Amoeba Robots

Locomotion in complex terrain can be influenced not only by environment but also by interactions with other individuals within the environment. We investigate the collective dynamics arising from confined active three-link particles (with two degrees of freedom). We refer to these as "smart" particles, or smarticles, because they are equipped with a microphone and can execute programmed behaviors based on the frequencies of external audio commands. The smarticles are confined inside an unanchored ring, which, due to the motion of the smarticles, can exhibit Brownian motion [30]. Smarticles are programmed to execute a square gait, which does not produce net translation or rotation in space for an individual; however, when confined together, the interactions between particles can generate net translations and rotations of the boundary in the lab frame. In contrast to previous studies of the collective motion of swarm systems, which have focused on movement of the collective based on the movement and activity of individuals [11], we are interested in the collective motion arising from the presence of immobile individuals, essentially creating robots made of robots.

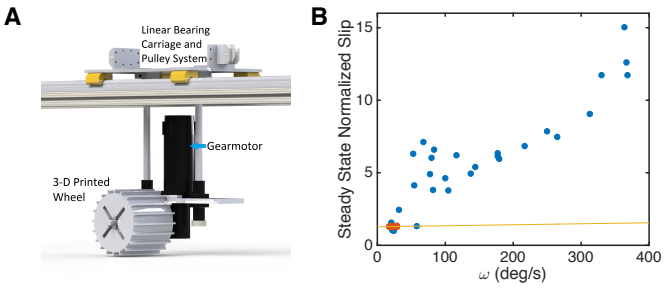


Fig. 14: Wheel translation in granular materials. (A) Digital rendering of the experimental single-wheeled locomotion test rig. With a pulley system adding counterweights, different effective gravities acting on the wheel can be achieved. Wheels can also be swapped out using the same gearmotor to test scaling relations between similar wheels of differing dimensions. (B) Red dots represent the steady state normalized slip for our collaborators performing a similar experiment over a slowly rotating limit, Blue dots represent the steady state normalized slip for a wider domain of rotation speeds in our setup. RFT suggests that slip should remain around 1.2, close to the yellow line which shows the nearly flat linear fit to the red dots. However, as rotation speed increases, normalized slip overshoots RFT's prediction.

We have found that we can indeed move between Brownian motion to directed motion of the system by introducing asymmetries into the system, such as placement of particles, which act differently than the rest of the collective. This may provide insight into how to induce and control a transition between a system exhibiting Brownian motion and one undergoing directed motion through small, potentially-minimally-invasive techniques. One particular method to do so is shown in the figure above where in experiment we were able to generate directed motion by the placement of a single inactive particle into the system.

I. Geometric mechanics

Geometric mechanics [31, 32] is a powerful framework relating position and orientation changes of a body to internal shape changes, and enables the prediction of effective movement and comparison between swimming in granular media and in viscous fluids [14], and advances now enable application to large self-deformations [13]. Previous work demonstrated the power of the combination of geometric mechanics and RFT through the study of a three-link robot with two active, controllable degrees of freedom (DoF) swimming in granular media [14].We have extended these tools to describe the locomotion of animals and robots with more than two DoF, showing that the cyclic self-deformations of sandfish lizards [20], shovel nosed snakes [33], sidewinder rattlesnakes [5, 22], salamanders and mudskipper fish [26] that make continuous or intermittent contacts with granular substrates can be represented as a linear combination of two shape bases functionsthe shape configuration space can be described with two variables (Fig. 16). With these approximations, we begin to gain insight into optimal self-deformation (templates

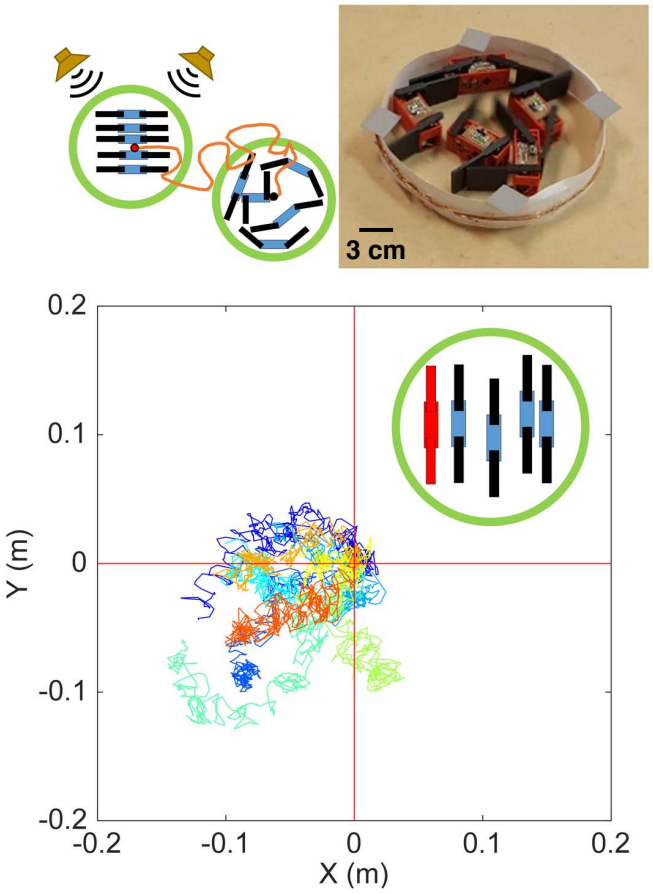


Fig. 15: Collective-motion-induced locomotion. (Top left) Experimental setup schematic for the "Amoebot." (Top right) Image of the amoebot. (Bottom) Time-varying position of the robots. Introducing a single inactive unit induces a directed motion in the robot.

in the language of [10]) of high-DoF locomotors in granular media).

IV. Summary

This work developed principles to characterize and model the dynamics of locomotion on complex flowable terrain. The progression of this research initiated with homogeneous and flowable terrain, and advanced to discrete obstacles and complex interactions with granular substrates. Using robots inspired by organisms from snakes to mudskippers to bipeds to amoebas, we show both how to negotiate complex terrain and how the interaction with substrates can yield complex behaviors. Phenomena such as wave-like diffraction patterns and relativistic motion emerged on the scale of 0.1 m to 1 m robots as a consequence of their interaction with heterogeneous and deformable substrates. This ongoing endeavor continues to build a general scientific framework [3] for enabling motion in natural environments.

References

[1] Jeffrey Aguilar and Daniel I Goldman. Robophysical study of jumping dynamics on granular media. *Nature Physics*, 2015.

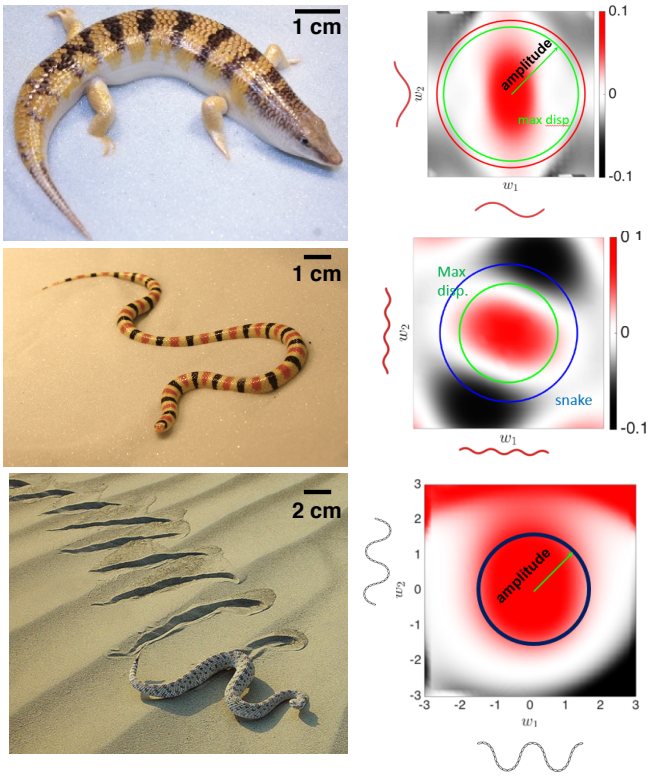


Fig. 16: Geometric mechanics predictions for animals with many degrees of freedom. Left column, from top to bottom, the sandfish skink, the Mojave shovel-nosed snake, and the sidewinder rattlesnake. These organisms use body bends to move on and within loose granular media. Right column, the geometric mechanics calculations for the organisms on the left. This is a new formulation of geometric mechanics which successfully predicts the kinematics of these many degrees-of-freedom animals.

[2] Jeffrey Aguilar, Alex Lesov, Kurt Wiesenfeld, and Daniel I Goldman. Lift-off dynamics in a simple jumping robot. *Physical review letters*, 109(17):174301, 2012.

[3] Jeffrey Aguilar, Tingnan Zhang, Feifei Qian, Mark Kingsbury, Benjamin McInroe, Nicole Mazouchova, Chen Li, Ryan Maladen, Chaohui Gong, Matt Travers, et al. A review on locomotion robophysics: the study of movement at the intersection of robotics, soft matter and dynamical systems. *Reports on Progress in Physics*, 79(11):110001, 2016.

[4] Aaron D Ames, Ramanarayan Vasudevan, and Ruzena Bajcsy. Human-data based cost of bipedal robotic walking. In *Proceedings of the 14th international conference on Hybrid systems: computation and control*, pages 153–162. ACM, 2011.

[5] Henry C Astley, Chaohui Gong, Jin Dai, Matthew Travers, Miguel M Serrano, Patricio A Vela, Howie Choset, Joseph R Mendelson, David L Hu, and Daniel I Goldman. Modulation of orthogonal body waves enables high maneuverability in sidewinding locomotion. *Proceedings of the National Academy of Sciences*, 112(19):6200–6205, 2015.

[6] Mieczyslaw Gregory Bekker. *Off-the-road locomotion: research and development in terramechanics*. University of Michigan Press Ann Arbor, 1960.

[7] Eric Brown, Nicholas Rodenberg, John Amend, Annan Mozeika, Erik Steltz, Mitchell R Zakin, Hod Lipson, and Heinrich M Jaeger. Universal robotic gripper based on the jamming of granular material. *Proceedings of the National Academy of Sciences*, 107(44):18809–18814, 2010.

[8] Aguilar J.J. Goldman D.I. Ames A.D. Vela P.A. Chang A., Hubicki C.M. Learning to jump in granular media: Unifying optimal control synthesis with gaussian process-based regression. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.

[9] Miguel Schneider Fontan and Maja J. Mataric. A study of territoriality: The role of critical mass in adaptive task division. In *In From Animals to Animats 4: Proceedings of the Fourth International Conference of Simulation of Adaptive Behavior*, pages 553–561. MIT Press, 1996.

[10] Robert J Full and Daniel E Koditschek. Templates and anchors: neuromechanical hypotheses of legged locomotion on land. *Journal of experimental biology*, 202(23):3325–3332, 1999.

[11] L. Giomi, N. Hawley-Weld, and L. Mahadevan. Swarming, swirling and stasis in sequestered bristle-bots. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 469(2151), 2013.

[12] Daniel I Goldman. Colloquium: Biophysical principles of undulatory self-propulsion in granular media. *Reviews of Modern Physics*, 86(3):943, 2014.

[13] RL Hatton and H Choset. Nonconservativity and noncommutativity in locomotion. *The European Physical Journal Special Topics*, 224(17-18):3141–3174, 2015.

[14] Ross L. Hatton, Yang Ding, Howie Choset, and Daniel I. Goldman. Geometric visualization of self-propulsion in a complex medium. *Phys. Rev. Lett.*, 110:078101, Feb 2013.

[15] Goldman D.I. Ames A.D. Hubicki C.M., Aguilar J.J. Tractable terrain-aware motion planning on granular media: An impulsive jumping study. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, pages 3887–3892, 2016.

[16] Shuuji Kajita, Fumio Kanehiro, Kenji Kaneko, Kiyoshi Fujiwara, Kensuke Harada, Kazuhito Yokoi, and Hirohisa Hirukawa. Biped walking pattern generation by using preview control of zero-moment point. In *Robotics and Automation, 2003. Proceedings. ICRA'03. IEEE International Conference on*, volume 2, pages 1620–1626. IEEE, 2003.

[17] Hiroaki Katsuragi and Douglas J. Durian. Unified force law for granular impact cratering. *Nature Physics*, 3(6):420–423, April 2007.

[18] Scott D. Kelly and Richard M. Murray. Geometric phases and robotic locomotion. *Journal of Robotic Systems*, 12(6):417–431, 1995.

[19] Chen Li, Tingnan Zhang, and Daniel I Goldman. A terradynamics of legged locomotion on granular media. *Science*, 339(6126):1408–1412, 2013.

[20] Ryan D Maladen, Yang Ding, Chen Li, and Daniel I Goldman. Undulatory swimming in sand: subsurface locomotion of the sandfish lizard. *science*, 325(5938):314–318, 2009.

[21] Ryan D Maladen, Yang Ding, Paul B Umbanhowar, and Daniel I Goldman. Undulatory swimming in sand: experimental and simulation studies of a robotic sandfish. *The International Journal of Robotics Research*, 30(7):793–805, 2011.

[22] Hamidreza Marvi, Chaohui Gong, Nick Gravish, Henry Astley, Matthew Travers, Ross L Hatton, Joseph R Mendelson, Howie Choset, David L Hu, and Daniel I Goldman. Sidewinding with minimal slip: Snake and robot ascent of sandy slopes. *Science*, 346(6206):224–229, 2014.

[23] N. Mazouchova, N. Gravish, A. Savu, and D.I. Goldman. Utilization of granular solidification during terrestrial locomotion of hatchling sea turtles. *Biology Letters*, 6(3):398, 2010.

[24] Nicole Mazouchova, Paul B Umbanhowar, and Daniel I Goldman. Flipper-driven terrestrial locomotion of a sea turtle-inspired robot. *Bioinspiration & Biomimetics*, 8(2):026007, 2013.

[25] B. McInroe, H. C. Astley, C. Gong, S. M. Kawano, P. E. Schiebel, J. M. Rieser, H. Choset, R. W. Blob, and D. I. Goldman. Tail use improves performance on soft substrates in models of early vertebrate land locomotors. *Science*, 353(6295):154–158, July 2016.

[26] Benjamin McInroe, Henry C Astley, Chaohui Gong, Sandy M Kawano, Perrin E Schiebel, Jennifer M Rieser, Howie Choset, Richard W Blob, and Daniel I Goldman. Tail use improves performance on soft substrates in models of early vertebrate land locomotors. *Science*, 353(6295):154–158, 2016.

[27] Feifei Qian and Daniel Goldman. Anticipatory control using substrate manipulation enables trajectory control of legged locomotion on heterogeneous granular media. In *SPIE Defense+ Security*, pages 94671U–94671U. International Society for Optics and Photonics, 2015.

[28] Feifei Qian and Daniel I Goldman. The dynamics of legged locomotion in heterogeneous terrain: universality in scattering and sensitivity to initial conditions. In *Robotics: Science and Systems*, 2015.

[29] Feifei Qian, Tingnan Zhang, Kevin Daffon, and Daniel I. Goldman. An automated system for systematic testing of locomotion on heterogeneous granular media. In *International Conference on Climbing and Walking Robots (CLAWAR)*, 2013.

[30] Michael J. Saxton. *Modeling 2D and 3D Diffusion*, pages 295–321. Humana Press, Totowa, NJ, 2007.

[31] Alfred Shapere and Frank Wilczek. Self-propulsion at low reynolds number. *Physical Review Letters*, 58(20):2051, 1987.

[32] Alfred Shapere and Frank Wilczek. Efficiencies of self-propulsion at low reynolds number. *Journal of fluid mechanics*, 198:587–599, 1989.

[33] Sarah S Sharpe, Stephan A Koehler, Robyn M Kuckuk, Miguel Serrano, Patricio A Vela, Joseph Mendelson, and Daniel I Goldman. Locomotor benefits of being a slender and slick sand swimmer. *Journal of Experimental Biology*, 218(3):440–450, 2015.

[34] J. Slonaker, D. Carrington Motley, C. Senatore, K. Iagnemma, and K. Kamrin. Geometrically general scaling relations for locomotion on granular beds. *ArXiv e-prints*, April 2016.

[35] Eric R Westervelt, Jessy W Grizzle, and Daniel E Koditschek. Hybrid zero dynamics of planar biped walkers. *IEEE transactions on automatic control*, 48(1):42–56, 2003.

[36] Jo Yung Wong. *Terramechanics and off-road vehicle engineering: terrain behaviour, off-road vehicle performance and design*. Butterworth-heinemann, 2009.

[37] Tingnan Zhang and Daniel I Goldman. The effectiveness of resistive force theory in granular locomotion a. *Physics of Fluids*, 26(10):101308, 2014.

[38] Esben H. stergaard, Gaurav S. Sukhatme, and Maja J. Mataric. Emergent bucket brigading - a simple mechanism for improving performance in multi-robot constrained-space foraging tasks. In *In Autonomous Agents*, pages 2219–2223, 2001.

3. Sensing, Perception, and Processing (SPP) Thrust

Thrust Lead: Kamal Sarabandi, University of Michigan
ARL Technical Lead: William Nothwang, Sensors and Electron Devices Directorate

3.1 Vision for Sensing, Perception, and Processing

To enable MAST, it is imperative for a microsystem, or system of microsystems, to be considered autonomous in that it must be able to understand and reason about its world. To accomplish this, at a minimum, a microsystem must be able to 1) measure itself and surroundings (e.g., sense), 2) reason about the sensed data (e.g., perceive), and 3) have the substrate to enable 1 and 2 (e.g., processing). Specifically, sensing refers to hardware required for monitoring the state of the platform, navigation, and ISR. Perception refers to interpretation of sensor data, and hence deals with algorithms and their embodiment in software. Processing here refers to computational hardware; while that is needed for many functions in MAST systems, perception is the biggest driver of processing requirements in MAST. SPP therefore encompass a closely related set of issues that we address as one of our cross-cutting thrusts. One salient feature of the sensors and algorithms developed for MAST platforms is the stringent requirement on SWaP. While commercial off-the-shelf (COTS) can and do supply many needs, such SWaP constraints often require very specific and intelligent solutions outside the existing approaches and COTS components. New hardware, software, and algorithms, as well as interfaces in between, have been developed to address the specific requirement of MAST platforms and their missions. Power, size, and weight improvements of 10 times to 100 times are required to meet the requirements for MAST robotic platforms. This necessitates exploitation of new physics and pushing the envelope of technology to achieve higher speed, density, frequency, and so on for different sensors considered. In addition, new ideas related to distributed sensing and sensor fusion had to be considered to accomplish what seemed impossible at the beginning of the program. It is also important to emphasize that the elements of SPP presented in this report are also closely related to elements of the other thrusts.

3.2 Technical Challenges

Based on the scenario-specific objectives described earlier, SPP scenario-driven capability requirements have been identified in state estimation, goal perception, obstacle detection and terrain recognition, mapping, ISR, and processing. In general, these can be grouped into Sensing and Estimation, 3-D Perception and

Mapping, and ISR. It is important to understand that environmental conditions and SWaP constraints have a major bearing on the capability requirements within these three areas. For example, MAST missions may take place in lit, dark, or optically obscured environments and in still or gusty air. Further, a representative weight budget for all sensing and processing hardware, including associated energy storage, ranges from a few grams to a few tens of grams; mission durations ranging from 10 to 30 min; and platform sizes from 10 to 30 cm long. With these bounding conditions, the required capabilities in each area that were addressed include the following:

1) **Sensing and Estimation**
   - Individual localization in absolute terms and relative to nearby objects.
   - Characterization of speed, acceleration, and angular acceleration
   - Team localization, knowing the relative positions of each team member.

2) **3-D Perception and Mapping**
   - Recognizing and tracking specified destinations as the robot approaches.
   - Omnidirectional obstacle detection and avoidance for air vehicles.
   - Detecting obstacles ahead of and to the sides of ground vehicles, as well as perceiving terrain properties as needed to adapt locomotion behavior (e.g., gait parameters) to the terrain.
   - 2- and 2.5-D mapping of single and multiple floors of intact buildings.
   - 3-D mapping of collapsed buildings and underground structures (caves, sewer systems).
   - Mapping complex outdoor/indoor/underground environments.

3) **ISR**
   - Detecting and tracking people and vehicles using their IR signatures
   - Optical sensor with high-field of view (FOV), low-light, and low distortion
   - Providing adequate computational throughput for all needed autonomous navigation and ISR functions, including low-level vision processing

3.3 Thrust Organization

To address the challenging needs of MAST scenario-specific objectives and platform requirements research activities in the SPP thrust is organized within the following three subthrusts:

1) Sensing and Estimation: Projects in this subthrust address the required capabilities for *state estimation* and *goal perception*.

2) 3-D Perception and Mapping: Projects in this subthrust address the required capabilities for *obstacle detection and terrain recognition* as well as *mapping.*

3) ISR: Projects in this subthrust are mission-oriented sensors and address the required capability for detection and assessment of threats.

3.4 Technical Challenges and Approach

3.4.1 Sensing and Estimation

Many of the technical challenges in state estimation (position, velocity, and attitude) are a consequence of the small size of MAST platforms, which mandates very small hardware elements for the solution. This often requires nonstandard solutions and excludes many of the existing commercial products. Additional difficulties come from the need to operate out of sight from the operator (NLOS), in GPS-denied environments with poor radio propagation or other types of RF navigation aids, in lit or dark conditions, and in the presence of gusts. Gusts, in particular, and high velocity aerial platforms, in general, require low latency and high accuracy state estimation. Since MAST scenarios involve teams of small platforms, it is also challenging to know the relative state of all members of the team with sufficient precision and low latency for effective mission execution. For maneuvering in tight spaces, state relative to other nearby objects can be more important than state relative to absolute reference frames. The tasks in this subthrust are summarized in Table 3-1.

Table 3-1 SPP technical challenges: Sensing and Estimation

Technical challenge	Performance objectives	Research tasks
High-performance sensing of acceleration and rotation	10 times better performance than COTS, 10 times lower power consumption	SPP-16-1.1, SPP-16-1.2
Fusion of stereo and optical flow sensors	Real-time obstacle avoidance for flying vehicles	SPP-16-1.3
Gust detection, airflow sensing	Real-time detection for inner-loop gust rejection control strategies. Optimized airfoil shape and motion for improved lift/drag characteristics.	SPP-16-1.4

Our approach draws on system dynamics identification work in the MCE thrust that helps to identify sensor response rates required for control of small platforms in disturbances. In the SPP thrust, we have conducted simulation studies and field

experiments to explore the use of state estimation algorithms to characterize state estimation performance as a function of the performance of individual sensors and the types of sensors composing the overall state estimation sensor suite (SPP-16-1.3). This provides guidance to select the types of sensor hardware to develop and establish performance goals for those efforts. Guided by the simulation and experimental results, we also have developed real-time algorithms for low latency state estimation, fusing visual, inertial, and altitude sensors (SPP-16-1.3). In particular, we sought to enable self-contained, onboard state estimation with low latency, low computation, high accuracy, and bandwidth, and a minimal sensor suite, to replace current solutions that use external position tracking sensors, like the GPS and motion capture systems, or large onboard sensors that are infeasible for small platforms, like the Hokuyo LADAR. Position knowledge can also be produced as a byproduct of mapping, though with greater latency and computational cost; we discuss this next.

To enhance the sensors available to MAST-scale platforms, we conduct research on biologically inspired concepts for hair-like accelerometers (SPP-16-1.1) and ocelli-like optical rotation sensors (SPP-16-1.2). These provide low-noise, high-bandwidth, accurate sensors that are superior to the COTS MEMS-based sensors that are now available. In particular, each is based on array technology: this provides many spatial estimates that can be averaged to reduce noise, be heterogeneous to increase bandwidth and noise rejection; or otherwise be processed with a higher update rate than is possible with a single noisy sensor that is averaged in time. Task SPP-16-1.4 is continuing efforts for arrayed, hair-like airflow sensors that will benefit the control algorithms, but also provide feedback as to flow over the control surfaces of airborne MAST platforms to solve the gust detection problem.

3.4.2 3-D Perception and Mapping

In NLOS, GPS-denied situations with highly constrained communication between robots and users, it can be difficult for a user to specify where a robot is supposed to go and for the robot to maintain knowledge of the location of the goal as the robot moves. The severe SWaP constraints of MAST robots makes this even harder by limiting their onboard perceptual capabilities. Goals can be represented many ways, from specifications as simple as direction vectors to follow, to 3-D coordinates in a given reference frame, to objects or physical structures that users can name and robots can recognize, to general directives to explore. The challenge was to develop a suite of goal perception capabilities that maximize mission utility while remaining within plausible SWaP constraints.

To facilitate development and testing of reactive and deliberative control algorithms on small crawlers, we integrated very simple goal perception sensors on small crawling robots, in the form of photodiodes and antennas for homing on light and basic visual target tracking capabilities to crawlers. We have been working on 3 approaches to recognizing objects or physical structures as goals: 1) a landing site recognition algorithm for aircraft that uses a form of visual structure-from-motion to discriminate roof-like structures from the ground, 2) a vision algorithm to recognize rectangular structures corresponding to doors or windows for building ingress, and 3) an algorithm that uses monocular vision together with a single-axis LADAR to detect doors inside a building. In the CNC thrust, we also have developed cooperative goal perception methods where an aerial vehicle provides situational awareness to help guide a small ground vehicle to a goal.

A great challenge for obstacle detection is getting sensors small enough. This is especially true for air vehicles, which have the tightest mass constraints and ideally would have some obstacle detection capability in all directions. It is important to be able to conduct missions in darkness, which places additional demands on sensors for obstacle detection. For ground vehicles, dramatic differences in mobility performance can sometimes be achieved by adapting locomotion parameters to terrain characteristics; the challenges here are how to recognize the need to adapt and how to know what new behavior to employ.

For obstacle detection, COTS LADARs are currently used on surrogate vehicles for obstacle detection and mapping; however, these sensors weigh 160 g and are not directly amenable to miniaturization down to a few grams. As such, significant effort has been going into developing a very small, 230–245 GHz radar with electronic beam steering for forward-looking operation, as well as research on low-power processing for optical flow (SPP-16-1.2) and other optical sensors (SPP-16-3.2) that will stimulate further development of optical and optical flow-based approaches to obstacle detection in future years. Besides obstacle detection, the radar work also supports mapping and could also support state estimation, while operating without the need for sunlight and in the presence of some atmospheric obscuration. The radar system, the culmination of 10 years of component, phenomenology and systems research, in the coming year will complete the fabrication process for the radar antenna and start integration with the high-frequency amplifiers, frequency multipliers, and filters. The waveform generation hardware is being addressed by a high-performance 20-GHz modulator that was developed in task SPP-16-2.2. This uses a mostly digital technology to reduce the required power and also make it more agile. This technology, while matured for this MAST application, has led to other research efforts to produce a single-chip

40-GHz radar that promises to be part of a variety of military and consumer applications.

Challenges for mapping are similar to those for obstacle detection, including needs for very small sensors, 3-D sensor coverage, and operation in all lighting conditions. In addition, mapping is comparatively expensive computationally, so challenges exist to reducing the computational complexity of mapping algorithms and/or developing compact, low-power processors for them. The tasks in this subthrust are summarized in Table 3-2.

Table 3-2 SPP technical challenges: 3-D Perception and Mapping

Technical challenge	Performance objectives	Research tasks
Forward-looking obstacle detection including operability in the dark and obscured atmospheres	235-GHz, electronically beam-steered radar; 50°, 30° azimuth/elevation scan, 2° resolution in azimuth; 5g; 3×4×0.1 cm antenna	SPP-16-2.1, SPP-16-2.2
Aerial flight of MAST-scale platforms	3-D state estimation and detection of obstacles for maneuvering in 3-D using vision from monocular camera	SPP-16-1.3
Stealthy mapping	Through-wall imaging	SPP-16-2.3

Research on sensors for state estimation, obstacle detection, and seeing in the dark also serves mapping (SPP-16-2.1). Work on algorithms for mapping straddles the SPP and CNC thrusts: while some of it pushes the state of the art in perception (SPP), much of it involves cooperative, multi-robot operation (CNC). This work also bears on state estimation for individuals and teams, albeit with more computational complexity and longer latency than the methods discussed earlier. Work in past years focused on 2-D mapping with scanning LADAR and monocular imagery.

To explore the possibility of imaging the interior of buildings and rooms without the occupants being alerted, Task SPP-16-2.3 is developing technology to use multiple robots with low-frequency radars to image through walls. This uses a cooperative multi-robot strategy to gather enough data quickly to make the imaging timely. Task SPP-16-2.4 is exploring the integration of several of these sensors on actual platforms to encourage development of new robot control strategies, including sensor fusion approaches. Sensor integration is a critical aspect of the overall MAST system, enabling low SWaP. SPP-16-2.5 is developing the systems level integration of sensors and behaviors developed in some of the other MAST tasks. In particular, this task is working with the teams from SPP-16-1.2 (optical flow sensors for motion sensing), SPP-16-1.4 (hair-based airflow sensing), and SPP-16-3.2 (stereo vision disparity accelerator). Each of these sensors/devices is

being integrated with MAST-like robots in both hardware and software to explore the synergies available to the robot with such a sensor complement onboard.

3.4.3 ISR

Two past efforts in ISR have involved research into devices to detect airborne traces of things of great interest for surveillance: nuclear radiation and chemical contaminants. In each case, scale-down of current technology required development of quite new and different technologies. For nuclear radiation detection as a high-voltage gap was used to generate a current in response to a single nuclear particle. Multiple miniature detectors were stacked to provide high detection rates with low power, weight, and volume. For detection of chemical traces in the atmosphere, a small gas chromatograph was created, requiring small pumps, separation columns, heaters, and chemical detectors, all with an eye to making the system small and low power yet with the detection ability of a larger unit.

This program currently focuses on ISR with imaging, for example to detect people. Many of the other sensors in subthrusts 1 and 2 can also be used for ISR. Again, SWaP constraints are paramount challenges. The need to operate in lit or dark environments creates an extra challenge for image-based ISR when the total mass budget for electronics may be a few tens of grams or less. Also, imagery must be interpreted onboard or by a user before it provides a useful ISR result, which is challenged by low onboard computing power and low communication throughput to the user. The tasks in this subthrust are summarized in Table 3-3.

Table 3-3 SPP technical challenges: ISR

Technical challenge	Performance objectives	Research tasks
See in the dark and detect warm objects at long range	NEDT < 5 mK, 0.05 g, 10 fps, 8 mW, 1024×768	SPP-16-3.1
Fast mapping and object recognition	Real-time 3-D mapping, threat detection	SPP-16-3.2

The research on an uncooled IR detector begun in FY11 uses a novel thermal IR detector array using pixels based on Q-amplified MEMS resonators (SPP-16-3.1). The sensor includes a reference device for self-calibration and an array of detectors with detector pitch of only 7 μm using resonant piezo and pyroelectric materials. If this research is successful, by FY17 it will produce a 1024 × 768 IR detector array with a pixel pitch of 7 μm, power dissipation of less than 8 mW, mass of 55 mg (with the optics the mass can go up to few grams), and a noise equivalent delta temperature (NEDT) of 5 mK, thereby enabling thermal imaging on MAST-scale platforms.

Another task, SPP-16-3.2, is developing efficient, ultra-low power implementations of structure from motion (SfM) algorithms, leveraging high-resolution images to create 3-D representations of the robot's environment, and bio-inspired small target motion detection (STMD) to provide low computational effort tracking of objects, with emphasis on humans and small enemy platforms.

3.5 Significant Accomplishments

Making a set of sensors that is relevant and useful for MAST-scale robots is a significant challenge. This thrust has made significant accomplishments toward that goal. In particular, a while new class of sensors/actuators was created using hydraulic amplification. The sensors from this research were used for airflow sensing as well as linear acceleration. Used in groups, it enabled low-power, low-noise yet very fast sensing. A sensor for airborne chemicals was also fabricated and demonstrated that enables highly sensitive and highly specific sensing of a variety of species relevant to the military, such as explosives. This sensor was designed to be chip-sized, with a very low power requirement, enabling it to be used along with conventional sensors in unconventional scenarios, such as tracing chemicals to their source, or "smelling" chemical residues on people.[1–8] For the notional flapping platform that is still in the future, we have devised a wing material that is a very lightweight, high-efficiency solar cell array. For the first time, flexible solar panels with thickness values of few microns transferred on Mylar and Kapton sheets were demonstrated. This enables the batteries to recharge when the platform is outdoors in the sun, lengthening the mission lifetime and enabling high-power applications to execute. Solar to electrical power efficiency of more than 20% were demonstrated. To achieve multifunctionality, traces of the solar cell array with proper RF traps are used as VHF and UHF antenna on the wings of flappers.[9–14] Radio technology has also been improved in complex environments through the use of a small, lightweight, battery-operated radio repeater that can be strategically placed by the robots during their surveillance of an area.[15–22] This can be used to improve connectivity between the robots in the group, but also, in conjunction with a few added sensors, the capability of leave-behind sensors for continuous surveillance. Transmission of video at low radio frequencies has also been enabled using advanced signal processing, multiple channels, and high-efficiency amplifiers. This enables video transmission in complex environments where the multipath issues are prohibitive for higher frequencies.

Further accomplishments include an insect-inspired self-motion sensor based on research by another researcher in a different thrust area. This vision-based motion system uses an array of motion sensors along with hardware weighting and summation to produce signals that are related to a variety of different self-motions

of the platform. While not an inertial motion sensor, it provides a useful alternative that is both low power and high accuracy. Some sensors depend on a significant amount of processing before their data become useful to the robot. Several of these have been targeted by this thrust to develop a fast, low-power hardware solution. Specifically, subthreshold very large scale integration (VLSI) technology are substantially advanced to enable low-power processing and data storage. For example, the stereo vision scenario. Our thrust has developed a low-power, high-throughput stereo vision processor that provides the depth for every pixel in 2 high definition (HD) video streams, at 30 fps. Current software techniques were implemented in low-power hardware for a more than 100-times speedup. This can then be used as an input to the simultaneous localization and mapping (SLAM) processor that most robots depend on for navigation these days. A related vision sensor that was developed is a low-power IR imager, with specialized calibration circuitry that requires no cooling, as well as very low-power readout circuitry. A sensor for obscured scenarios, or sensing in the dark, is the 240-GHz beamsteering radar that is about the size and weight of a credit card. This provides a 30-fps slice of the surroundings, providing a range/angle representation that can be used for collision avoidance and mapping, similar to a much heavier and power-hungry LIDAR with significantly higher range unhampered by smoke, dust, and inclement weather conditions. For situational awareness and threat assessment, an ensemble of coherent ultra-wideband receivers and transmitters mounted on ground robots are used to realize a 2-D synthetic aperture radar for the first time to enable seeing through walls with 360° FOV to provide a very-high-resolution radar images of building environments.

Robust vison-based algorithms using multi-sensory multi-robots to create a map of the environment and determine the presence of human and animals have also been demonstrated. The algorithms used feature elements (lines, points, and planes) and a model-based estimator in addition to using temporal context to reduce complexity. Other vision based algorithms for localization and control are also developed. Stereo vision and optical flow has also been demonstrated for path planning collision avoidance.

3.6 Challenges for the Future

One of the challenges for the future involves devising sensors that are more integrated with their function, generally providing for low-power, high-speed hardware processing of the signals. This provides information instead of raw sensor data. Progress has begun in this area as evidenced by the stereo camera processor discussed earlier. We can envision smarter sensors, as well, that provide state information and can be throttled as needed for better acuity, and provide power

savings otherwise. The sensor fusion idea continues to be an area where great progress can be made, aiding the robot to better understand their environment and their situation. Again, hardware implementations of key algorithms will make it possible to have op tempos that are more relevant to actual mission needs. Exploiting the new paradigm of arrays of low-power, low-quality sensors to produce high-quality, high-speed results is an area that still has many opportunities for application in a great variety of sensing modalities.

3.7 SPP Task Summary

CENTER LEAD: Kamal Sarabandi, University of Michigan			
ARL TECH LEAD: Brian Sadler			
Project #	**Project Title**	**PI**	**ORG**
SPP-16-1	Sensing and Estimation	Sarabandi	Michigan
SPP-16-1.1	Biomimetic Super HAIR Inertial Sensor	Najafi	Michigan
SPP-16-1.2	Optical Flow Motion Sensing	Yoon	Michigan
SPP-16-1.3	Stereo-Optical Flow Fusion in Inverse Depth Space for Obstacle Avoidance	Matthies/ Kuwata / Brockers	JPL
SPP-16-1.4	High-Density Biomimetic Microhydraulic Air Flow Sensor Array	Najafi	Michigan
SPP-16-2	Localization, Mapping and Sensory Fusion	Sarabandi	Michigan
SPP-16-2.1	Advanced Submillimeter Wave Radar for Sensing & Navigation on MAST Platforms	East Sarabandi	Michigan
SPP-16-2.3	Through-the-Wall Radar Imaging Using an Ad hoc Ensemble of MAST	Sarabandi	Michigan
SPP-16-2.4	Air Ground Mapping of Feature Rich 3-D Indoor Environments with RGBD Sensors	Kumar / Hsieh / Taylor	Pennsylvania
SPP-16-2.5	Sensor Exploitation in Airborne Mobile Robots	Josyula	Bowie
SPP-16-3	ISR Payloads	Rais-Zadeh	Michigan
SPP-16-3.1	Low-Power Readout for Resonant Sensors	Rais-Zadeh	Michigan
SPP-16-3.2	Energy Efficient Approaches to Enable Advanced Computer Vision for MAST Platforms	Sylvester / Blaauw	Michigan

3.8 References

1. Agah, M., Potkay, J. A., Lambertus, G., Sacks, R., and Wise, K. D. (2005). High-performance temperature-programmed microfabricated gas chromatography columns. *Journal of Microelectromechanical Systems*, *14*(5), 1039-1050.

2. Potkay, J. A., Lambertus, G. R., Sacks, R. D., and Wise, K. D. (2007). A low-power pressure-and temperature-programmable micro gas chromatography column. *Journal of Microelectromechanical Systems*, *16*(5), 1071-1079.

3. Kim, S. J., Reidy, S. M., Block, B. P., Wise, K. D., Zellers, E. T., and Kurabayashi, K. (2010). Microfabricated thermal modulator for comprehensive two-dimensional micro gas chromatography: design, thermal modeling, and preliminary testing. *Lab on a Chip*, *10*(13), 1647-1654.

4. Liu, J., Gupta, N. K., Wise, K. D., Gianchandani, Y. B., and Fan, X. (2011). Demonstration of motionless Knudsen pump based micro-gas chromatography featuring micro-fabricated columns and on-column detectors. *Lab on a Chip*, *11*(20), 3487-3492.

5. Kim, H., Steinecker, W. H., Reidy, S., Lambertus, G. R., Astle, A. A., Najafi, K., and Wise, K. D. (2007, June). A micropump-driven high-speed MEMS gas chromatography system. In *Solid-State Sensors, Actuators and Microsystems Conference, 2007. TRANSDUCERS 2007. International* (pp. 1505-1508).

6. Lu, C. J., Steinecker, W. H., Tian, W. C., Oborny, M. C., Nichols, J. M., Agah, M., and Wise, K. D. (2005). First-generation hybrid MEMS gas chromatograph. *Lab on a Chip*, *5*(10), 1123-1131.

7. Zellers, E. T., Reidy, S., Veeneman, R. A., Gordenker, R., Steinecker, W. H., Lambertus, G. R., and Avery, C. (2007, June). An integrated micro-analytical system for complex vapor mixtures. In *Solid-State Sensors, Actuators and Microsystems Conference, 2007. Transducers 2007. International* (pp. 1491-1496). IEEE.

8. Wise, K. D., Sacks, R., Beach, K. T., Potkay, J. A., and Agah, M. (2005). *U.S. Patent No. 6,838,640*. Washington, DC: U.S. Patent and Trademark Office.

9. Shiu, K. T., Zimmerman, J., Wang, H., and Forrest, S. R. (2009). Ultrathin film, high specific power InP solar cells on flexible plastic substrates. *Applied Physics Letters*, *95*(22), 223503.

10. Wei, G., Xiao, X., Wang, S., Zimmerman, J. D., Sun, K., Diev, V. V., and Forrest, S. R. (2011). Arylamine-based squaraine donors for use in organic solar cells. *Nano letters*, *11*(10), 4261-4264.

11. Che, X., Xiao, X., Zimmerman, J. D., Fan, D., and Forrest, S. R. (2014). High-Efficiency, Vacuum-Deposited, Small-Molecule Organic Tandem and Triple-Junction Photovoltaic Cells. *Advanced Energy Materials*, *4*(18).

12. Lee, K., Shiu, K. T., Zimmerman, J. D., Renshaw, C. K., and Forrest, S. R. (2010). Multiple growths of epitaxial lift-off solar cells from a single InP substrate. *Applied Physics Letters*, *97*(10), 101107.

13. Lee, K., Zimmerman, J. D., Hughes, T. W., and Forrest, S. R. (2014). Non-Destructive Wafer Recycling for Low-Cost Thin-Film Flexible Optoelectronics. *Advanced Functional Materials*, *24*(27), 4284-4291.

14. Lee, K., Zimmerman, J. D., Zhang, Y., and Forrest, S. R. (2012, June). Epitaxial lift-off of GaAs thin-film solar cells followed by substrate reuse. In *Photovoltaic Specialists Conference (PVSC), 2012 38th IEEE* (pp. 001698-001700). IEEE.

15. Sarabandi, K., and Song, Y. J. (2011). Subwavelength radio repeater system utilizing miniaturized antennas and metamaterial channel isolator. *IEEE transactions on antennas and propagation*, *59*(7), 2683-2690.

16. Oh, J., Thiel, M., and Sarabandi, K. (2014). Wave-propagation management in indoor environments using micro-radio-repeater systems. *IEEE Antennas and Propagation Magazine*, *56*(2), 76-88.

17. Song, Y. J., and Sarabandi, K. (2012). Miniaturized radio repeater for enhanced wireless connectivity of ad-hoc networks. *IEEE Transactions on Antennas and Propagation*, *60*(8), 3913-3920.

18. Oh, J., Thiel, M., Hong, W., and Sarabandi, K. (2009, June). Indoor wave propagation measurements and modeling for evaluation of coverage enhancement using a repeater system. In *Antennas and Propagation Society International Symposium, 2009. APSURSI'09. IEEE* (pp. 1-4). IEEE.

19. Song, Y. J., and Sarabandi, K. (2010, July). Small radio repeater system for enhancement of wireless connectivity. In *Antennas and Propagation Society International Symposium (APSURSI), 2010 IEEE* (pp. 1-4). IEEE.

20. Song, Y. J., and Sarabandi, K. (2012). Equivalent circuit model for metamaterial-based electromagnetic band-gap isolator. *IEEE Antennas and Wireless Propagation Letters*, *11*, 1366-1369.

21. Song, Y. J., and Sarabandi, K. (2014). A simultaneous dual-channel micro-radio-repeater for ad-hoc wireless communication. *IEEE Transactions on Antennas and Propagation*, *62*(6), 3378-3383.

22. Amjadi, S. M., and Sarabandi, K. (2016). Mutual coupling mitigation in broadband multiple-antenna communication systems using feedforward technique. *IEEE Transactions on Antennas and Propagation*, *64*(5), 1642-1652.

Single-Chip Optic Flow Sensor

SPP-17-1.2

Kyueok Lee and Euisik Yoon
Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, MI 48109-2122, USA
esyoon@umich.edu

Joseph Conroy, PhD
Army Research Laboratory
2800 Power Mill Rd.
Adelphi, MD 20783, USA
joseph.k.conroy3.civ@mail.mil

Integrated image sensors are essential for navigation and autonomous flight of ground robots and micro air vehicles (MAVs). Estimation of optic flow velocity is of great importance and very much needed for local path finding and collision avoidance in the operation of MAST platforms. Implementation of single-chip optic flow sensors will allow compact, low-payload, low-power deployment of optical navigation modules for a wide field of view by using an adaptable configuration in system integration. Using the bio-inspired time-stamped algorithm, we have achieved a wide optic flow range of up to 2 rad/sec at a rate of 120 frames/sec for reliable operation under extreme conditions. Also, we implemented the low-power object detection in mixed-mode circuits for potential applications of rescue and surveillance. The task of the current fiscal year has been focused on the implementation of reliable optic flow sensing in a single-chip by integrating all the timing and control circuits embedded with the core pixel array along with the object detection unit.

***Index Terms*—Optic-Flow Sensor, Autonomous Flight, Micro Air Vehicles.**

I. Introduction

In the previous years, we have developed a couple of prototype optic flow sensors. We validated our custom neuromorphic algorithm to mimic insect vision in the mixed-mode circuits using 0.18μm CMOS technology at extremely low power consumption of less than 30 μW (2-D optic flow core circuit) or 243.3pJ/pixel. This bio-inspired time-stamp optic flow estimation algorithm allows effective tracking of moving features at 120 fps with the maximum measurable optic flow range of 1.96 rad/sec. However, the previous prototype devices require external control units including a clock generator and bias circuits and interface logics. This makes the assembled module huge, not suitable to be deployed for a MAST platform.

In the final year, we focused on integrating all the control and peripheral circuits in the same chip, realizing a stand-alone single-chip optic flow sensor. We demonstrate the stand-alone single-chip operation through a miniaturized module for MAST Capstone demonstration. The implemented optic flow sensor chip can be easily utilized and deployed in the MAST platform through the SPI bus interfacing with a host microcontroller. Furthermore, fully-integrated single chip optic flow sensors can be mounted on a hemispherical module for a wide field of view in the deployment of various MAST platforms in the future. The package and lens amounting should be supported by additional funding for the full deployment of the prototype optic flow module in a MAST platform.

II. Related works

Traditionally, the air-vehicle systems used separate image sensors and digital processors/controllers, however this approach results in high payload and huge power consumption. For the conventional digital optic flow estimation system and algorithms, such as Lucas-and-Kanade, huge numerical calculations are required. Typically, a substantial digital hardware (CPU and/or FPGA) should be allocated and dedicated for numerical number crunching, resulting in large power consumption [1], [2], [3]. As an alternative approach to reduce power consumption, a bio-inspired elementary motion detector (EMD) have been introduced and implemented in analog VLSI circuits for autonomous navigation [4]–[8]. However, pure analog signal processing is easily susceptible to temperature changes and process variations. It is also known that scaling of analog circuits is difficult because it is not trivial to transport the design from one process technology to another. In the proposed prototype device, we have implemented a time-stamp-based optic flow algorithm, which is modified from the conventional EMD algorithm to give an optimum partitioning of hardware blocks in analog and digital domains as well as assign adequate allocation of pixel-level, column-parallel, and chip-level processing.

III. approach

The sensor operates in two different modes: optic flow (OF) estimation mode and object detection (OD) mode. The estimation of optic flow is a task that measures the speed of moving features. This task involves the measurement of time and a distance for a moving object (or feature) across pixels. In the OF estimation mode, the sensor generates a 64×64 16b OF data at 30 fps using bio-inspired time-stamped algorithm. The pixel consists of sampling capacitor, gain capacitor and OTA (Operational Transconductance Amplifier) to monitor temporal contrast changes. Previous and current frames are compared using the sampling capacitance and temporal

differences are amplified by PGA (Programmable Gain Amplifier). A 1b-comparator detects a moving feature when the temporal contrast is beyond threshold. The column-parallel 1b-feature information updates recent time-stamp information and the bio-inspired 2D optic-flow core estimates 16b OF data. The generated OF data is serialized, compressed and sent to the host.

The sensor switches its mode from OF extraction to OD once (or twice) in every second (programmable). For the OD mode, we implemented the histogram-of-oriented-gradients (HOG) to provide robust operation for object detection against illumination variation and various backgrounds. Typically, the HOG feature and support vector machine (SVM) require a complicated calculation, huge memory and high-resolution images. In this work we implemented the gradient orientation from 2b spatial difference images (SDI) and cell-based classification to save both memory area and power. The processing of OD is as follows: (1) a cell histogram generator assembles the magnitude of gradients corresponding to its 9 bins in an 8x8 pixel subarray; (2) the accumulated histograms are normalized with 4 neighbor cells; and (3) a linear SVM classifies the object by 9b normalized HOG features using cell-based pipeline operation. The SVM classification consists of 128 processing elements (PE) to support cell-based pipeline operation. The search window sweeps all windows. When the object is detected, additional searches are conducted for 4 neighboring windows. If the number of positive results among 4 windows is bigger than threshold, the system provides a positive result. Using this search scheme, the false negative values can be decreased to < 6%.

IV. RESULTS

In this year, we concentrate our efforts to implement and demonstrate the stand-alone single-chip module operation. Fig. 1, 2, and 3 show the captured optic flow (OF) images and object detection result from the fabricated device. To demonstrate the performance and feasibility in actual MAV scenarios, we tested 3 cases: a spiral-moving-up object, a rotating fan at 65 rpm, and a walking person. In Fig. 1, black arrows show the estimated OF of the spiral-moving-up object from the fabricated chip, while the red arrow indicates the actual trajectory of the object. The measured optic flows match well with the trajectory in 3D motion. Fig. 2 shows the OF estimation from a rotating fan (65rpm). This result explicitly shows the evaluation of accuracy in optic flow estimation in terms of optic flow directions and values.

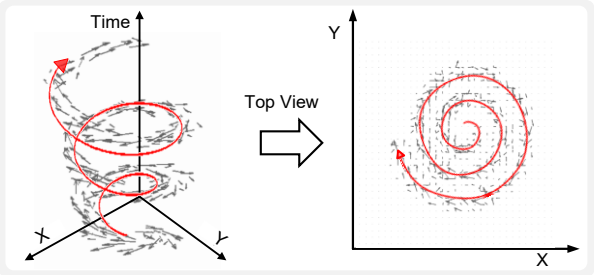


Fig. 1. Optic flows measured for a spiral-moving-up object @ 30fps

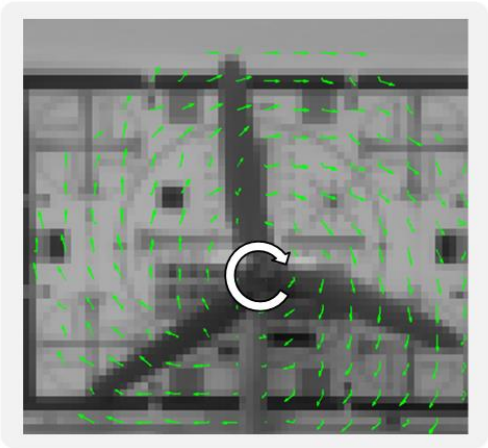


Fig. 2. Rotating fan: @ 65prm, captured @ 60fps

In Fig. 3, the sensor illustrates the actual scenario of the deployment of the sensor in an MAV. The prototype sensor generates optic flows during the OF extraction mode, showing in green arrows. Every once in a while (every 30th frame or every one second), the sensor switches to the feature extraction mode and classify the object from the scene. Yellow rectangles show the classified object (in the case the human object). The location of a walking person is indicated inside the yellow box. The performance of the sensor is summarized in Table 1.

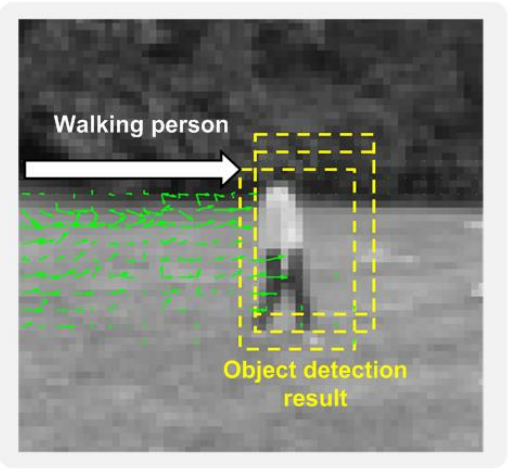


Fig. 3. Optic flow estimation (green arrows) of a walking person and object detection (yellow rectangles) when operated in feature extraction mode. Optic flows are captured @ 30 fps and object detection/classification is conducted @ 1 fps.

TABLE I. PERFORMANCE TABLE

Process	0.18 μm 1P4M CMOS		
Chip size	6.20 x 4.00 mm^2		
Pixel size	31 x 31 μm^2		
Pixel array	64 x 64 (Optic flow extraction mode) / 256 x 256 (Object recognition mode)		
Fill factor	19%		
Maximum optic flow	1.96 rad/sec @ 120fps, FOV 60°		
2-D optic flow estimation (@ 30fps)	29.94 μW (244 pJ/pixel)	Power(pixel, 3.3V)	0.58 μW
		Power(Analog, 1.8V)	5.33 μW
		Power(2-D optic flow core, 1.8V)	24.03 μW
Object recognition (@ 30fps)	2.18 mW (1.11 nJ/pixel)	Power(pixel, 3.3V)	1.94 μW
		Power(Analog, 1.8V)	261.4 μW
		Power(HOG & SVM core, 1.2V)	1.92 mW
Hybrid operation mode Optic flow estimation mode @ 29fps Object recognition mode @ 1fps	101.61 μW (272.49 pJ/pixel)		
Detection rate of object recognition	84 %		

To verify the performance of the embedded feature extraction and object classification core, we carried out object detection using 200 pedestrian images from INRIA dataset [10] and achieved a detection rate of 84 %. For optic flow estimation, we have accomplished a FOM (figure of merit) of 272.49pJ/pixel in the hybrid mode operation. A prototype chip is fabricated using a 0.18µm CMOS process. A chip micrograph is shown in Fig. 4.

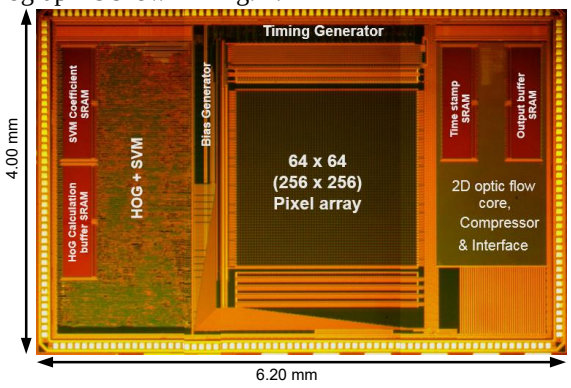


Fig. 4. Chip photograph

V. CONCLUSIONS AND FURTHER WORK

We have successfully demonstrated the stand-alone single-chip miniaturized module at the 2017 MAST capstone meeting. Fig. 5 show the miniaturized module, which provides a low payload and a small form factor. Eventually, the sensors will be packaged in a hemispherical assembly module for artificial compound eyes, but this will require additional support for design and implementation of small lens mounts and flexible PCBs.

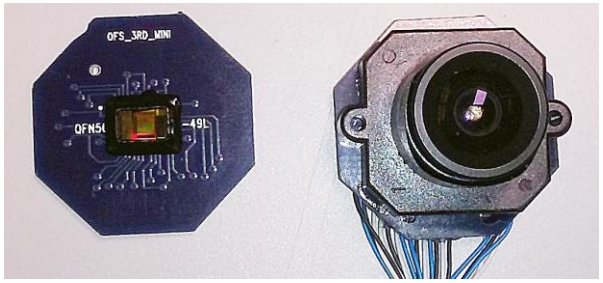


Fig. 5. Sensor assembled in a minimal PCB module

REFERENCES

[1] J. S. Humbert and A. M. Hyslop, "Bioinspired Visuomotor Convergence," IEEE Trans. Robot., vol. 26, no. 1, pp. 121–130, 2010.

[2] J. Conroy, G. Gremillion, B. Ranganathan, and J. S. Humbert, "Implementation of wide-field integration of optic flow for autonomous quadrotor navigation," Auton. Robots, vol. 27, no. 3, pp. 189–198, Aug. 2009.

[3] V. Mahalingam, K. Bhattacharya, N. Ranganathan, H. Chakravarthula, R. R. Murphy, and K. S. Pratt, "A VLSI Architecture and Algorithm for Lucas–Kanade-Based optical flow computation," vol. 18, no. 1, pp. 29–38, 2010.

[4] J. Krammer and C. Koch, "Pulse-based analog VLSI velocity sensors," IEEE Trans. Circuits Syst. II Analog Digit. Signal Process., vol. 44, no. 2, pp. 86–101, 1997.

[5] R. R. Harrison, "A biologically inspired analog IC for visual collision detection," IEEE Trans. Circuits Syst. I Regul. Pap., vol. 52, no. 11, pp. 2308–2318, Nov. 2005.

[6] A. A. Stocker, "Analog Integrated 2-D Optical Flow Sensor," Analog Integr. Circuits Signal Process., vol. 46, no. 2, pp. 121–138, Dec. 2005.

[7] R. S. A. Brinkworth, P. A. Shoemaker, and D. C. O'Carroll, "Characterization of a neuromorphic motion detection chip based on insect visual system," in 2009 International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP), pp. 289–294, 2009.

[8] P. Xu, J. S. Humbert, and P. Abshire, "Analog VLSI Implementation of Wide-field Integration Methods," J. Intell. Robot. Syst., vol. 64, no. 3–4, pp. 465–487, Feb. 2011.

[9] S. Park, J. Cho, K. Lee, and E. Yoon, "243.3pJ/pixel bio-inspired time-stamp-based 2D optic flow sensor for artificial compound eyes," in 2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2014, pp. 126–127.

[10] INRIA Person Dataset, [online] http://pascal.inrialpes.fr/data/human/, accessed Nov. 10, 2015.

Metrics

Notes: *Please list metrics for the reporting quarter only and only Publications that are accepted and searchable on web.*

Awards and/or Recognitions this year

Title

Refereed/Journal Papers accepted this year

Title

Conference Presentations/proceedings presented this year

Title
"A 272.49 pJ/pixel CMOS image sensor with embedded object detection and bio-inspired 2D optic flow generation for nano-air-vehicle navigation," IEEE Symposium on VLSI Circuits, pp. C294-C295, 2017

Joint Papers with ARL presented or accepted this year– duplication with journals/ conferences expected.

Title

Invention Disclosures for this year

Title

Patent Applications Filed this year

Title

Patents Awarded this year

Title

PhDs Award this year:

Name

Masters Awarded this year

Name

Students supported fully or partially (by name)

MS: *Name(s)*
PhD: *Hyunsoo Song, Kyuseok Lee*
Post Docs: *Sung-Yun Park,*
Master Degree: *Name(s*
PhDs Awarded: *Name(s*

MAST CTA Member Rotations into ARL

Name(s

ARL Researcher Working at Member Sites

Name(s

Workshops and Seminars conducted:

Title

MAST Quarterly Metrics Summary Table

MAST ANNUAL REPORT METRICS TABLE (The data entered into this table should be for the reporting quarter only)	
Awards (Miscellaneous Recognitions and Awards)	0
Refereed/Journal papers	0
Conference presentations/proceedings	1
Joint Papers with ARL	0
Invention disclosures	0
Patent applications filed	0
Patents awarded	0
PhDs Award	0
Masters Awarded	0
Post Docs Supported	1
Master Degree Students Supported	0
Undergraduate Students Supported	0
CTA Member rotations into ARL	0
Visiting Researchers from ARL to the Member Sites	0
Employees hired by ARL from CTA students supported	0

Vision-based Obstacle Avoidance for MAVs using an Egocylindrical Depth Map

SPP-16-17-1.3

L. Matthies, R. Brockers, A. Fragoso, B. Rothrock, and C. Lee
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA
brockers@nasa.jpl.gov

***Abstract*—Obstacle avoidance is an essential capability for micro air vehicles. Prior approaches have mainly been either purely reactive, mapping low-level visual features directly to headings, or deliberative methods that use onboard 3-D sensors to create a 3-D, voxel-based world model, then generate 3-D trajectories and check them for potential collisions with the world model. Onboard 3-D sensor suites have had limited fields of view. We use forward-looking stereo vision and lateral structure from motion to give a very wide horizontal and vertical field of regard. We fuse depth maps from these sources in a novel robot-centered, cylindrical, inverse range map we call an egocylinder. Temporal fusion is performed in two fashions: (1) by generalizeing the stereo disparity map to a Gaussian Mixture Model (GMM) with multiple depth hypotheses per pixel, which are updated over time given the temporal sequence of depth maps, and (2) by temporal propagation of depth data on the egocylinder, to maintain a persistent, robot-centered, 360° representation beyond the instantaneous field of view of the sensors. Configuration space expansion directly on the egocylinder gives a very compact representation of visible freespace. This supports very efficient motion planning and collision-checking with better performance guarantees than standard reactive methods. We show the feasibility of this approach experimentally in a challenging outdoor environment.**

***Index Terms*— Miniature, Micro Air Vehicle, Autonomous Navigation, MAST-scale**

I. Introduction

A. Background and Technical Challenges

Micro air vehicles (MAVs) require the ability to perceive and avoid obstacles while flying at any velocity indoors, outdoors, and underground, ultimately in all ambient lighting conditions from full sunlight to completely dark. This requires a perception system with a wide horizontal and vertical field of regard, high angular and range resolution, and the ability to work in any illumination condition, all while maintaining a very low size, weight, and power budget (SWaP). This is an extremely challenging sensor selection problem, and no single type of sensor meets this challenge. World model data structures for obstacle avoidance must have adequate resolution to represent relatively small obstacles at short range, adequate azimuthal extent to enable reasoning about sharp turns at low speed, and sufficient downrange extent to represent distant obstacles during high speed flight, while supporting motion planning, collision checking, and sensor and temporal fusion with a relatively low memory footprint and computing load. Simultaneously satisfying these requirements is also extremely challenging.

In prior work under the MAST program [1], we addressed this problem by drawing inspiration from the vision system of barn owls, which have eyes with a large binocular field of view plus large, monocular, peripheral visual fields [2]. While it is uncertain exactly how owl vision works, such a camera arrangement in MAVs allows 3-D perception in the forward direction whether the aircraft is moving or stationary, allows 3-D perception of distant objects to the sides through monocular structure from motion, and provides wide overall horizontal and vertical fields of view. Progress in very high resolution CMOS imagers promises to enable relatively large maximum range and range resolution with small baselines, potentially with on-chip pixel binning for multi-resolution, region-of-interest readout to provide a very power-efficient approach to foveal processing and focus of attention [3]. We generalized prior work on motion planning in image space [4] by showing that a configuration space (c-space) type of expansion can be applied to stereo disparity images to compactly represent free space for MAVs and to enable very efficient collision checking of candidate trajectories projected into image space. Using a Closed-Loop RRT (CL-RRT) motion planner [5] that generated dynamically feasible trajectories on the fly, we demonstrated this approach with a quadrotor using stereo vision for obstacle avoidance in a small grove of trees.

In FY15-16, we generalized this work to include two additional cameras (one looking to each side), to perform monocular structure from motion with these cameras, and to extend the scene representation from just a forward-looking, planar image space into an egocentric cylindrical image space ("egocylinder") that enables fusion of depth data from the forward-looking stereo cameras and the peripheral view cameras. Due to high computational cost of the CL-RRT planner, we replaced it with a very simple planner that uses heuristics to choose a goal direction from the egocylinder and

to steer the vehicle toward the goal; this is updated on every perception and control cycle. The egocylinder-based perception pipeline has been tested on real image data, the reactive obstacle avoidance algorithm has been tested with a synthetic world simulator and with real data, and all components of the approach have been implemented in real-time in onboard processing architecture and flight tested in a grove of trees [6].

This approach addresses many of the technical challenges outlined above. A number of key problems remain, including coping with the very high dynamic range of intra-scene illumination (e.g. full sunlight and deep shadows in the same image) and inter-scene illumination (i.e. large illumination changes frame to frame) that is often experienced, coping with missing or wrong depth estimates from the vision algorithms, and inability of simple, reactive control laws for obstacle avoidance to guarantee collision free flight as the obstacle density and MAV velocity increase.

B. Research Objectives

Our objectives are to:

1. Improve the ability of this system to operate in scenes with a high dynamic range of illumination.
2. Improve the ability of the system to cope with missing and wrong depth data.
3. Develop a new computationally efficient motion planning algorithm formulated with the image space representation that ensures the generation of dynamically feasible, collision-free trajectories for highly cluttered scenes and high velocity flight.

II. Related Work

Pros and cons of various passive and active sensor options for MAV obstacle avoidance were discussed in [1, 7]; recent examples of MAV systems using multiple types of sensors are described in [8, 9]. Here we focus on methods that use vision.

Vision-based approaches break down according to how they do vision, scene representation, and planning and control. The main approaches to vision use optical flow, learning, and/or monocular or stereo depth perception. Optical flow methods typically design reactive control algorithms with optical flow input. Control algorithms for provably stable wall-following and corridor-following behavior have been developed this way [10]; however, navigation that requires a discrete choice among alternate directions requires higher-level perception or reasoning. Machine learning methods have also been used to map optical flow and other monocular visual features into reactive obstacle avoidance behaviors [11], but it is difficult to generalize this approach to work in a wide variety of conditions, so most work on MAV obstacle detection uses depth perception. Forward-looking monocular depth perception via structure from motion (SfM) has been used for MAVs [12, 13], but requires aircraft motion to measure depth and has poor depth perception near the focus of expansion. Stereo vision overcomes these limitations, works well in many outdoor settings in particular, and small, fast stereo implementations are progressing [14, 15, 16].

The predominant approach to scene representation has been to use 2-D or 3-D Cartesian probabilistic grid maps, which can be used with motion planning algorithms that vary from reactive to deliberative and from local to global [8, 16, 17, 18]. These methods are particularly useful for mapping, exploration, or obstacle avoidance in areas that require memory of previously examined avenues; however, they use a lot of storage and computation, and scaling to high speed flight requires multi-resolution grids.

For obstacle avoidance *per se*, less expensive representation and planning algorithms are possible. Often these representations are polar in nature, matching the polar angular resolution of the depth sensors [19, 20]. In [4], depth data from two onboard stereo pairs was fused in a cylindrical inverse range map centered on the vehicle. This work introduced c-space obstacle expansion of an image space depth map, though in a limited fashion based on an assumed ground plane. In [1], we generalized the c-space expansion to be based on the actual depth at each pixel and developed the first combination of an image space depth representation with a dynamics-aware motion planner. Feasible trajectories were generated in 3-D and projected into image space to do collision checking by testing for intersections with the c-space expanded depth map. This approach to obstacle representation and collision checking is fast and showed good potential in experiments; however, the field of view was limited and the CL-RRT motion planning algorithm was computationally expensive.

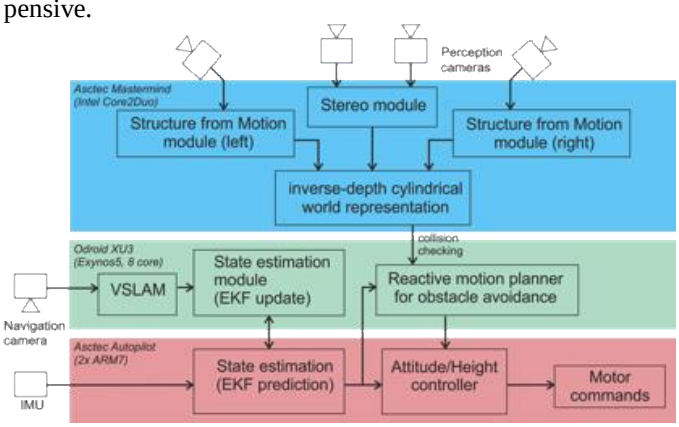


Figure 1. System architecture.

III. Approach

Figure 1 illustrates the sensing, processing, and algorithm architecture of our approach, which is implemented on an AscTec Pelican quadrotor (Figure 5). To minimize the sensor hardware, we augment forward-looking stereo by adding single cameras aimed 45° to each side, giving a total field of regard of about 180°. Stereo matching is done with local block matching for speed, which works adequately well in our test environments. To obtain depth perception with the side-looking cameras, we examined several options for using two-frame optical flow algorithms, as well as the LSD-SLAM incremental SfM algorithm [21]. LSD-SLAM constrains optical flow search to epipolar lines computed from estimated camera motion, and incrementally improves depth resolution by using

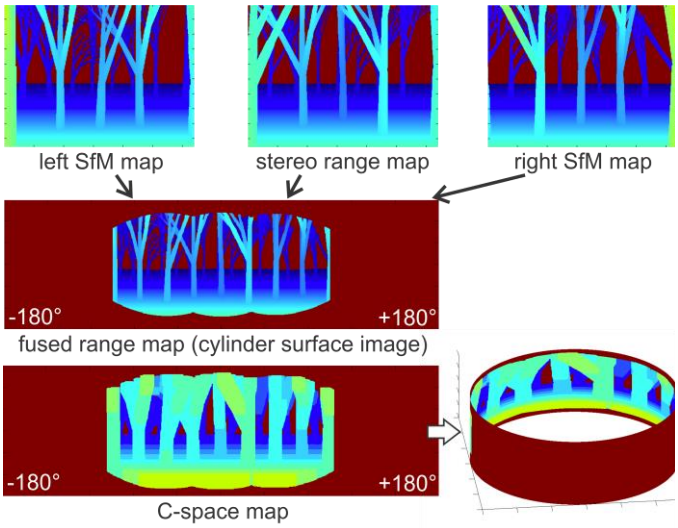


Figure 2. Schematic illustration of stereo and SfM depth maps fusion into the egocylinder representation, and c-space expansion of the egocylinder. Using inverse range, the expansion widens closer objects more than farther objects.

each new image to update depth estimates in keyframes. We found this to be much less noisy than unconstrained optical flow algorithms, so we use this approach. Since monocular SfM has an unobservable scale factor, we estimate scale by comparing SfM range measurements with range from stereo in the image overlap areas between the outer SfM cameras and the stereo cameras.

To provide a unified depth representation, we project the stereo and scaled SfM depth maps onto a cylinder centered on the aircraft, which is stored as a disparity map where each pixel records the inverse radial distance to the nearest object in the direction of the pixel (Figure 2). Currently the egocylinder has the orientation of the body frame, though it could be aligned with the gravity vector. C-space expansion is done on this disparity map similarly to [1], which essentially reduces the range at each pixel and widens objects in the depth map in proportion to the radius of the aircraft plus a safety margin. Using inverse range conveniently gives a finite representation (zero) to objects beyond the maximum range of the sensors. Quantized inverse range also matches the range uncertainty characteristics of vision-based depth estimation.

Traditional voxel-based 3-D world models perform temporal fusion by using successive frames of depth data to update statistics in each voxel that indicate whether the cell is empty or occupied. This involves 3-D ray-tracing of pixels in the depth data. We leverage work done under MAST task JX-16-17-1 on a new approach to temporal fusion directly in stereo disparity space, which generalizes the depth representation at each pixel to a Gaussian Mixture Model (GMM) that represents multiple depth hypotheses at each pixel and updates these over time [23]. This achieves the usual goals of temporal fusion of filling in missing data and reducing the frequency of gross errors; it also has potential to improve the precision of depth estimates in Kalman filter-like fashion, though we have not evaluated this performance to date. Currently, results from this temporal fusion step overwrite the corresponding portion of the egocylinder representation. Depth data on the egocylinder is propagated in time around the cylinder as the aircraft moves to maintain a persistent, robot-centered, 360° representation of the depth to the nearest surface in all directions around the vehicle. Our most recent work extends this method to provide a low-cost method for coping with moving objects in the scene [25].

The c-space expanded egocylinder allows the aircraft to be treated as a point for collision checking. It also reduces the dimensionality of search for a collision free trajectory from 3-D to 2-D, because we can search the 2-D space of visual directions for intermediate waypoints that are in front of the depth recorded in the egocylinder, generate a polynomial trajectory to the waypoint, and then check the trajectory for potential collisions by a z-buffer-like test to ensure that all points on the trajectory are also in front of the depths recorded in the eogcylinder. If a collision is found, the polynomial can be constrained to make a tighter turn to avoid the collision, or a new intermediate waypoint can be chosen in the egocylinder. Details are discussed in [26]. Once a collision-free, dynamically feasible trajectory is found, controls are generated by standard methods based on configuration flat dynamics.

Figure 1 shows how all parts of the algorithm map onto our three-level processor architecture. Images are processed at 384x240 pixel resolution. Stereo runs at 5 fps, LSD-SLAM at 10 fps, and the egocylinder and motion plan are updated at the stereo frame rate of 5 fps. Planning itself takes a few millisecond. Visual-inertial state estimation is done with a nadir-pointed camera using methods from [22].

Operating outdoors in areas with bright sunlight and deep shadow is difficult, because it creates very large intra-scene (within the same image) and inter-scene (between successive images) illumination variations that greatly exceed the linear dynamic range of available cameras. This has been especially problematic in experiments we have conducted in a grove of trees (see Section 4) using a nadir-pointed camera for state estimation. The most effective way to address this is to improve dynamic range at the sensor level. There are multiple potential ways to do this. Some approaches acquire multiple images separated in time and combine these in software; this is impractical on a moving robot. Another approach uses hardware design in the imager that provides a multi-linear, high dynamic range (HDR) exposure mode that approximates a logarithmic image response. This mode is implemented in the Matrix Vision mvBlueFOX-200w CMOS cameras we use and can extend the total dynamic range from 55 dB to 110 dB. These cameras have three linear segments in their photometric response function, where the slope and transition point of the second and third segments is controlled by a two sets of knee point parameters. Creating a good exposure for given scene conditions requires choosing the total exposure time and setting appropriate knee point parameters.



Figure 4. Non-HDR (left) and HDR (right) images in a forest scene. Large areas are saturated or under-exposed in the non-HDR image. The HDR image has a better distribution of intensity values, which leads to better performance of vision algorithms.

We have taken a first step toward exploiting this multi-linear HDR mode in the following camera initialization procedure, which is run once at the start of an experiment (Figure 4). First, we acquire a series of images while adjusting exposure time via gradient descent to push the average intensity of the image stream toward a target intensity in the middle of the pixel brightness range. Next, we fix the total exposure time while seeking the parameters of each knee point that maximize image entropy. In an iterative process, each knee point is set sequentially to maximize local entropy. This does not simultaneously optimize the setting of both knee points, but it avoids extra parameters and improves feature tracking performance. Once the exposure parameters are initialized, they are fixed for the duration of the flight, which has been adequate in our test conditions to date. Ideally, exposure should be optimized on every frame; however, our current optimization procedure is too slow for that and large changes of exposure have potential to require changes to feature tracking algorithms to maintain landmark tracking across exposure discontinuities. The latter issue is out of scope at present.

IV. Results

We have conducted low-speed (< 1 m/s) experimental trials in a grove of trees that provided a relatively high obstacle frequency (Figure 5). Flights totaled over 500 meters in aggregate length, during which 65 trees were encountered. This area had very difficult illumination conditions due to the combination of brightly sunlit and deeply shadowed areas in the same image. Figure 6 shows results of the vision pipeline at several points during such a run, as well as a 2-D map produced after the fact from data logs.

The saturated and underexposed areas of the images in Figure 4 illustrate the dynamic range problem with these illumination conditions. While the c-space expansion effectively fills in many areas that have missing data in depth maps from stereo and SfM, this forest environment was particularly challenging for the visual-inertial state estimation system. Therefore, we focused HDR experiments on the state estimation camera, where use of the HDR mode improved the average percentage of map landmarks that could be matched in each frame from 61% to 79%. Nevertheless, the floor of the forest had many very small, self-similar features, and doing state estimation with a nadir-pointed camera while low (< 2 m above the ground) in this environment still made state estimation by far the weakest link in the system.



Figure 5. Left: Grove of trees test area; Right: AscTec Pelican with 4 camera head.

The obstacle detection and avoidance portions of the architecture were very reliable in the performance evaluation experiments, which were analyzed quantitatively by noting the frequency and cause of any human intervention required to avoid a collision. These modules were responsible for only a single intervention event during the 521 meters recorded, which resulted in successful avoidance of 64 out of 65 trees for a success rate of 98%. The intervention was attributed to a missed detection in which the vehicle had drifted too close to an obstacle and could no longer detect it using stereo matching. LSD-SLAM failed to adequately track features about 25% of the time; with data logging, the LSD-SLAM frame rate dropped to about 8 Hz, which is too slow for this algorithm to be reliable. However, this did not impact overall obstacle avoidance performance, because the control policy of first turning the stereo cameras towards the flight direction, and incorporating a small amount of path hysteresis, provided a high degree of robustness to missed left or right camera LSD-SLAM depth maps, which were seamlessly reacquired beginning on the next frame.

V. Conclusions and Further Work

Using C-space expansion of image space depth maps for collision checking is a very new approach to obstacle avoidance for MAVs. In our experiments to date, obstacle avoidance has been quite successful; in 521 meters of flight in challenging conditions, only one intervention was needed in 65 encounters with obstacles, and no problems with false alarms in freespace were apparent. Offline evaluation has illustrated the benefits of GMM-based temporal fusion for reducing the occurrence rate of gross errors in depth data and filling in data that is missing in individual frames [23, 25, 26].

By far the biggest performance problem in this system is with visual state estimation. Using a nadir pointed camera while flying low (< 2 m above ground) in a scene with a very high dynamic range of illumination and many small, self-similar features (leaves) on the forest floor seems to be at the heart of the problem. We plan to address this in several ways in ongoing work, including using visual feature tracking in the forward and sideward looking cameras. Using tightly coupled visual-inertial state estimation algorithms, such as the Multi-State Constrained Kalman Filter (MSCKF), would provide more reliability, but require more computational power than the Odroid available in our current processing architecture.

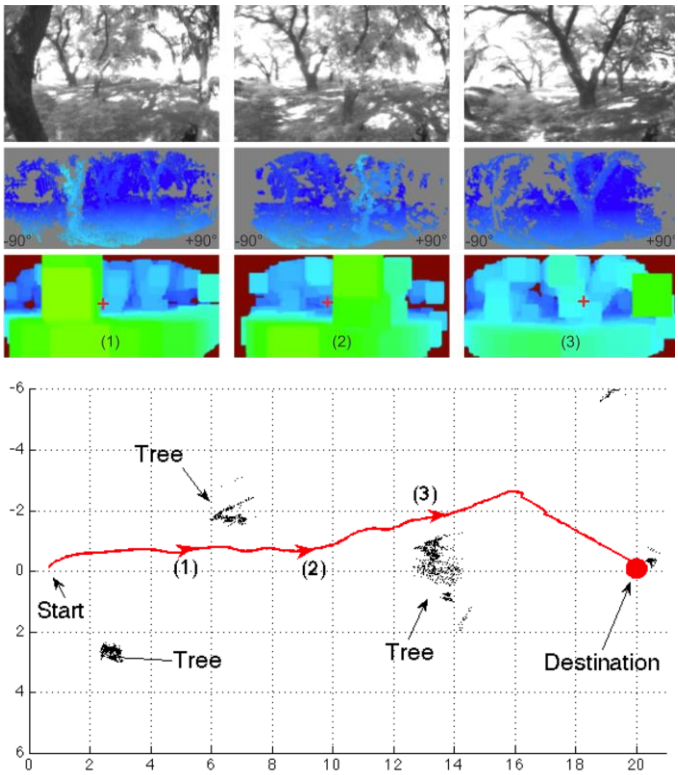


Figure 6. Results of a 20 m experimental flight through a grove of trees. Top: the results of the perception system for three different locations on the run, showing the left rectified stereo image, the fused egocylinder, and the c-space expanded egocylinder with selected direction of flight (red crosses). This only shows the central 180° of the egocylinder. Bottom: a top down 2-D plot of the trajectory and nearby obstacle pixels from the egocylinder over the whole run. Arrows and numbers on the trajectory show where the three images above were acquired. Vehicle speed was 1 m/s throughout.

The disparity-space planner extends the advantages of the c-space expansion method and egocylinder to the planning regime -- potential trajectories are selected and executed in a highly economical fashion by employing the same framework that allows the egocylinder to represent obstacles compactly and efficiently. Overall, this choice of representation demonstrates decreased planning latency and complexity compared to world-coordinate methods.

There is a close connection between vehicle velocity, uncertainty in the range data, and successful obstacle avoidance -- this has not emerged as an issue for the slow speeds of our experiments so far, but for reliable obstacle detection to scale to high speeds, this interplay will require further study. Several approaches may improve the maximum range and range resolution of the system to support higher velocities, including the use of higher resolution imagery with emerging ASIC processors for fast, low power stereo matching. Scenes involving moving obstacles will require extensions of both the perception and the planning elements of this system; initial steps in that direction for perception using our GMM temporal fusion framework for moving object detection are described in [25].

REFERENCES

[1] L. Matthies, R. Brockers, Y. Kuwata, and S. Weiss. "Stereo vision-based obstacle avoidance for micro air vehicles using disparity space," IEEE International Conference on Robotics and Automation (ICRA), 2014.

[2] G. R. Martin. "What is binocular vision for? A bird's eye view," Journal of Vision, 9(11):1-19, 2009.

[3] S. Kemeny, R. Panicacci, B. Pain, L. Matthies, and E. Fossum. "Multiresolution image sensor," IEEE Transactions on Circuits and Systems for Video Technology, 7(4), 1997.

[4] M. Otte, S. Richardson, J. Mulligan, and G. Grudic. "Path planning in image space for autonomous robot navigation in unstructured outdoor environments," Journal of Field Robotics, 26(2), 2009.

[5] Y. Kuwata, J. Teo, G. Fiore, S. Karaman, E. Frazzoli, J.P. How, "Real-time motion planning with applications to autonomous urban driving", *IEEE Transactions on Control Systems Technology*, 17(5), 2009.

[6] R. Brockers, A. Fragoso, B. Rothrock, C. Lee, and L. Matthies, "Vision-based obstacle avoidance for micro air vehicles using an egocylindrical depth map," International Symposium of Experimental Robotics, 2016.

[7] F. Kendoul, "Survey of advances in guidance, navigation, and control of unmanned rotorcraft systems," Journal of Field Robotics, 29(2):315-378, 2012.

[8] D. Droeschel, M. Nieuwenhuisen, M. Beul, D. Holz, J. Stuckler, and S. Behnke, "Multi-layered mapping and navigation for autonomous micro air vehicles," Journal of Field Robotics, 33(4):451-475, 2015.

[9] S. Nuske, S. Choudhury, S. Jain, A. Chambers, L. Yoder, S. Scherer, L. Chambelain, H. Cover, and S. Singh, "Autonomous exploration and motion planning for an unmanned aerial vehicle navigating rivers," Journal of Field Robotics, 32(8):1141-1162, 2015.

[10] J. Keshavan, G. Gremillion, H. Alvarez-Escobar, and J. S. Humbert, "Autonomous vision-based navigation of a quadrotor in corridor-like environments," International Journal of Micro Air Vehicles, 7(2):111-123, 2015.

[11] D. Dey and et al. Vision and learning for deliberative monocular cluttered flight," 10th Conference on Field and Service Robotics, 2015.

[12] H. Alvarez, L. M. Paz, J. Sturm, and D. Cremers, "Collision avoidance for quadrotors with a monocular camera," 14th International Symposium on Experimental Robotics, 2015.

[13] D. Daftry, D. Dey, H. Sandhawalia, S. Zeng, J. A. Bagnell, and M. Hebert, "Semidense visual odometry for monocular navigation in cluttered environments," IEEE International Conference on Robotics and Automation, Workshop on Recent Advances in Sensing and Actuation for Bioinspired Agile Flight, 2015.

[14] S. B. Goldberg and L. Matthies, "Stereo and IMU assisted visual odometry on an OMAP3530 for small robots," IEEE Conference on Computer Vision and Pattern Recognition, Workshop on Embedded Computer Vision, June 2011.

[15] M. Kuhn, S. Moser, O. Isler, F. K. Gurkaynak, A. Burg, N. Felber, H. Kaelin, and W. Fichtner, "Efficient ASIC implementation of a real-time depth mapping stereo vision system," IEEE 46th Midwest Symp. on Circuits and Systems, 2003.

[16] K. Schmid, P. Lutz, T. Tomic, E. Mair, and H. Hirschmuller, "Autonomous vision-based micro air vehicle for indoor and outdoor navigation," Journal of Field Robotics, 31(4):537-570, 2014.

[17] F. Fraundorfer, L. Heng, D. Honegger, G. H. Lee, L. Meier, P.

Tanskanen, and M. Pollefeys, "Vision-based autonomous mapping and exploration using a quadrotor MAV," IEEE/RSJ International Conference on Intelligent Robots and Systems, 2012.

[18] S. Shen, N. Michael, and V. Kumar, "3D indoor exploration with a computationally constrained MAV," Robotics: Science and Systems, 2011.

[19] M. Bajracharya, A. Howard, L. Matthies, B. Tang, and M. Turmon, "Autonomous on-road navigation with end-to-end learning for the LAGR program," Journal of Field Robotics, 26(1):3-25, 2009.

[20] H. Oleynikova, D. Honegger, and M. Pollefeys, "Reactive avoidance using embedded stereo vision for MAV flight," IEEE International Conference on Robotics and Automation, 2015.

[21] J. Engel, T. Schops, and D. Cremers, "LSD-SLAM: Large-scale direct monocular SLAM," European Conference on Computer Vision (ECCV), September 2014.

[22] S. Weiss, M. Achtelik, S. Lynen, M. Achtelik, L. Kneip, M. Chli, and R. Siegwart, "Monocular vision for long-term micro aerial vehicle state estimation: a compendium," Journal of Field Robotics, 30(5), 2013.

[23] C. Cigla, R. Brockers, and L. Matthies, "Gaussian mixture models for temporal depth fusion," Winter Conference on Applications of Computer Vision (WACV), 2017.

[24] J. Engel, J. Stuckler, and D. Cremers, "Large-scale direct SLAM with stereo cameras," IEEE/RSJ International Conference on Intelligent Robots and Systems, September 2015.

[25] C. Cigla, R. Brockers, and L. Matthies, "Image-based visual perception and representation for collision avoidance," CVPR Embedded Vision Workshop, June 2017.

[26] A. Fragoso, C. Cigla, R. Brockers, and L. Matthies, "Dynamically feasible motion planning for micro air vehicles using an egocylinder," 11th Conference on Field and Service Robotics, September 2017.

SPP-17-1-3 FINAL REPORT – METRICS

Conference Presentations/proceedings presented this quarter

- A. Fragoso, L. Matthies, and R. Murray, "A fast motion planning representation for configuration flat robots with applications to micro air vehicles," American Control Conference, 2017.
- A. Fragoso, C. Cigla, R. Brockers, and L. Matthies, "Dynamically feasible motion planning for micro air vehicles using an egocylinder," 11th Conference on Field and Service Robotics, 2017

Patent Applications Filed this year

Controlling unmanned aerial vehicles to avoid obstacle collision.

Students supported fully or partially

PhD: Tony Fragoso (Caltech)

Post-docs supported fully or partially

Cevahir Cigla

MAST FY16 ANNUAL REPORT METRICS TABLE	
Awards (Miscellaneous Recognitions and Awards)	0
Refereed/Journal papers	0
Conference presentations/proceedings	2
Joint Papers with ARL	0
Invention disclosures	0
Patent applications filed	0
Patents awarded	0
PhDs Award	0
Masters Awarded	0
Post Docs Supported	1
PhD Degree Students Supported	1
Master Degree Students Supported	0
Undergraduate Students Supported	0
CTA Member rotations into ARL	0
Visiting Researchers from ARL to the Member Sites	0
Employees hired by ARL from CTA students supported	0

A Micromachined Packaging with Incorporated RF-Choke for Integration of Active Chips at SubMillimeter-wave Frequencies – SPP-17-2.1

Armin Jam, *Student Member, IEEE,* Jack East, *Life Member, IEEE* and Kamal Sarabandi, *Fellow, IEEE*

***Abstract*— This paper presents the design, fabrication and evaluation of a novel packaging method for integration of active MMIC at millimeter-wave (MMW) and sub-MMW frequencies. This packaging approach is intended for integration with the MAST micromachined sub-MMW radar system operating at 240 GHz in order to complete the signal chain from the MMIC to the antennas. The concept is based on a peg and socket structure composed of RF transitions and biasing circuitry where the RF and DC pads on the MMIC chip mate with the on-wafer RF transitions and biasing pads, respectively, for a seamless packaging method. To prevent instability of the packaged active device, an ultra-wideband RF-choke is devised to provide better than 15 dB of RF isolation between the chip and the biasing circuitry from 30 GHz to 330 GHz. A prototype of the RF-choke along with the rest of the packaging block is fabricated using multi-layer silicon micromachining technology at Y-band (i.e. 220-325 GHz) and the performance is measured where close agreement of the measured and simulated results is shown. Additionally, the proposed integration method is used for packaging of an amplifier MMIC and a thru-line chip where excellent RF and biasing results are achieved.**

***Index Terms*— Integration, millimeter-wave, MMIC, radar, RF-choke, silicon micromachining, submillimeter-wave, terahertz.**

I. Introduction

THE next frontier in microwave engineering technology is in sub-millimeter-wave (sub-MMW) and terahertz (THz) band. The advent of solid state electronics has led into higher mobility of electrons in these devices which in turn has enabled realization of MMIC devices with F_{MAX} of beyond 1 THz [1], [2]. The development of such sub-MMW MMIC modules opens the door to whole new applications ranging from super-high resolution imaging systems with resolutions on the order of cm to 100 GB/s communication systems [3], [4].

Despite all these advancements, a key challenge still needs to be addressed in order to enable realization of such systems. This challenge pertains to the integration of the MMIC modules with other components such as high performance antennas, filters, couplers, etc. in these systems. The prevalent idea to address this issue is to monolithically integrate all modules including all passive components together with the active devices on the wafer the MMICs are being fabricated. However, considering the size of all passive and active modules and taking into account the high fabrication cost, this approach is cost limited especially in case where a relatively large aperture is required. Additionally, integrating all modules on a single chip often results into design tradeoffs between different modules due to the shared fabrication process.

To overcome this problem, a packaging method is required to enable reliable integration of the active MMIC with the rest of sub-MMW components. Most of these components are based on waveguide structures [5]-[7]. This is partly due to the superior loss performance of waveguides compared to other types of transmission lines. Another advantage of waveguides

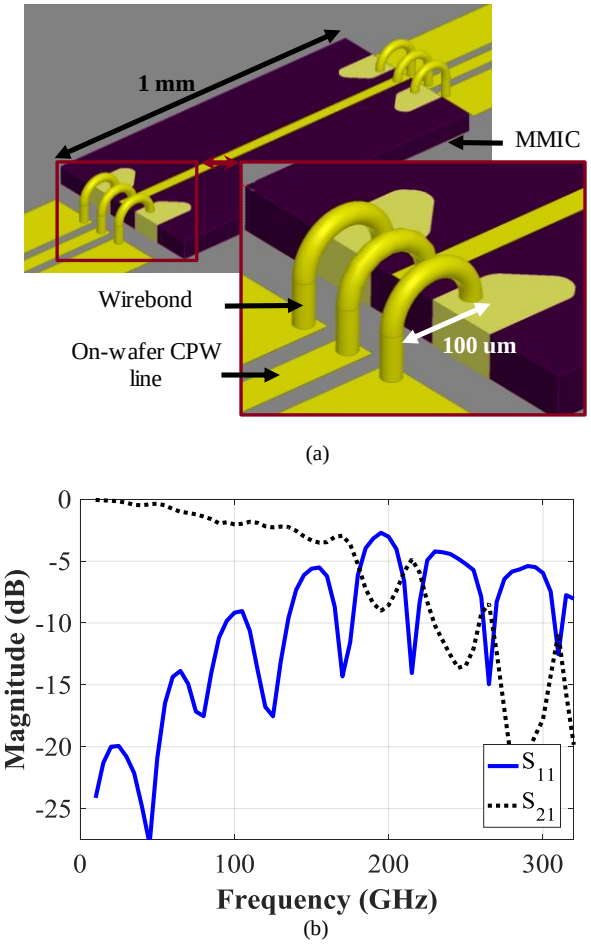


Fig. 1. (a) Diagram of a wire-bonded chip for integration of active MMIC with the rest of the system, and (b) Full-wave simulated *S*-parameters of the back-to-back wire-bonded thru-line chip showing deteriorated performance of packaging with the increase in the frequency.

is their compatibility with standard fabrication techniques at sub-MMW frequencies, such as silicon micromachining. Active MMIC modules, on the other hand, are most conveniently fabricated and integrated on planar transmission lines. To connect such devices to the rest of the system, appropriate transitions and packaging methods are required.

The dominant approach used at lower MMW frequencies is to use the wirebond technique to integrate MMICs with the on-wafer lines. However, the performance of this method is hampered at higher MMW frequencies as the effects of the parasitics associated with the wirebod are not negligible anymore (as shown in Fig. 1(b)). An alternative approach for integration of MMIC at lower MMW frequencies is to use the flip-chip technique, however, in this method the MMIC is exposed to high temperatures and pressures which is not desirable. A number of approaches for integrating MMIC modules with the rest of the system at the higher MMW frequencies have been reported in the literature [8], [9]. The principle of these methods is based on converting the coplanar or microstrip mode in the lines on the MMIC to the lines of the passive portion of the system using a probe transition that is monolithically integrated on the MMIC chip. Although good performances have been reported, these transitions have complex 3-D geometries which require complicated machining fabrication process with assembly of various parts, which is expensive.

In this work a novel chip integration and packaging methodology for sub-MMW applications is introduced. The integration method is based on a peg and socket approach with transitions from the line on the active MMIC chip to the fabricated lines on a packaging block. In order to enable scalability to the THz regime, the transition and the mounting block are designed to be compatible with silicon micromaching technology which offers a standardized fabrication process with very high accuracies. This results into a compact and cost effective integration and packaging solution in the sub-MMW to THz frequency band which can be standardized for reliable and repeatable integrations at such frequencies. In order to enable biasing of the active MMIC, a seamless biasing network consisting of an ultra-wideband RF-choke with a decade bandwidth (30-330 GHz) as well as the associated DC pads are devised as part of the integration block. The broad-band isolation is particularly important to prevent possible instability of the circuits, as the devices have higher gains at lower portion of the MMW frequencies. In order to validate the performance of the proposed integration methodology a prototype is designed and microfabricated at *Y*-band (220-325 GHz) and excellent measured performance is reported. The new integration method is readily compatible with MMIC and silicon micromachining technology and should find a wide range of applications in sub-MMW frequencies. Here, we present results on this integration and packaging methodology.

This paper is organized as follows. In Section II, the design of the integration method with the associated transitions and the RF-choke is discussed. Section III presents the microfabrication process of the proposed package and finally in Section IV the measurement results of packaged MMIC are presented.

II. Integration Structure and Design

A. Packaging Structure

Fig. 2 shows a diagram of the peg and socket chip integration configuration consisting of two bonded wafers that from here onwards are referred to as top and bottom wafers. The bottom wafer consists of an opening to insert and mount the MMIC chip. The size of the opening is chosen based on the size of the chip to allow accurate placement and alignment of the chip with the lines and pads on the wafer. The chip is supported from the top by two ledges on the top wafer over the opening in the bottom wafer. The chip is inserted through the opening from bottom and then supported by a piece of silicon block etched to the size of the hole. The top wafer consists of a membrane that RF and DC lines are patterned on

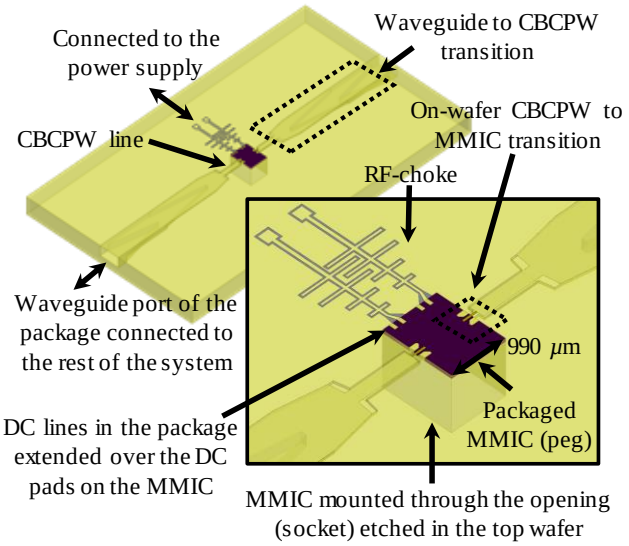


Fig. 2. Diagram of the packaging block with the associated transitions and the bias circuitry for integration of active MMIC in the micromachined radar system.

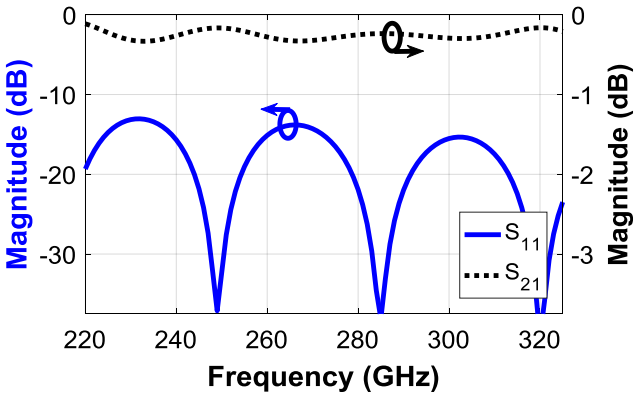


Fig. 3. Full-wave simulated *S*-parameters of the back-to-back transition from the on-wafer CBCPW line to the microstrip line on a delay line chip. The results only include the losses associated with the package.

it according to the position of the DC and RF pads on the MMIC chip. With the chip mounted in the opening, the traces on the membrane and the pads on the MMIC chip overlap and a physical contact is established between them. Also to facilitate DC contact between the two, many metallic cylinders with height of 5 μm and diameter of 10 μm are plated on the traces on the top wafer. The traces are next transitioned to RF cavity-backed coplanar waveguide (CBCPW) lines that extend to the next module in the system using a CBCPW to rectangular waveguide transition [10]. As will be shown this method enables effective integration of the chip with other micromachined system modules with high level of accuracy and is a strong candidate for a standardized methodology for chip packaging and integration at sub-MMW and THz frequencies as it does not require wire bond or soldering.

B. Transition Design

To validate the performance of the integration method a prototype at *Y*-band is designed and simulated using a full-wave solver (HFSS). In this design a 50 Ω CBCPW line is used to integrate different modules in the system. Although use of other type of lines is also possible, the CPW has the advantage of being readily compatible with most of MMIC applications. This is due to the fact that CPWs are the most widely used planar transmission lines in the MMIC chips due to their ease of fabrication and integration of components in series or shunt. The cavity-backed feature also overcomes the inherent high insertion loss of the conventional CPW lines due to the presence of the dielectric substrate. The removal of the substrate in the CBCPW design also eliminates the excitation of the substrate modes and also supports the propagation of the TEM mode in the line which makes it preferable for wideband applications. Fig. 3 illustrates the simulated RF performance of a back-to-back transition from the on-wafer CBCPW line to the microstrip line on a delay line chip. For a fair evaluation of the transition the insertion loss associated with the dielectric loss in the microstrip line and the conductor loss in the CBCPW line is eliminated. The result shows that the back-to-back transition has a reflection loss of better than 13 dB with an insertion loss of less than 0.4 dB over the desired bandwidth, i.e. 220-325 GHz.

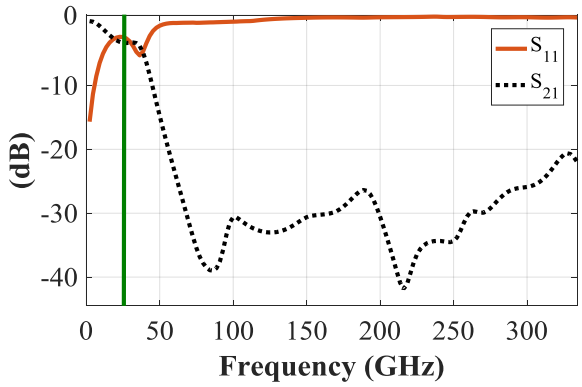
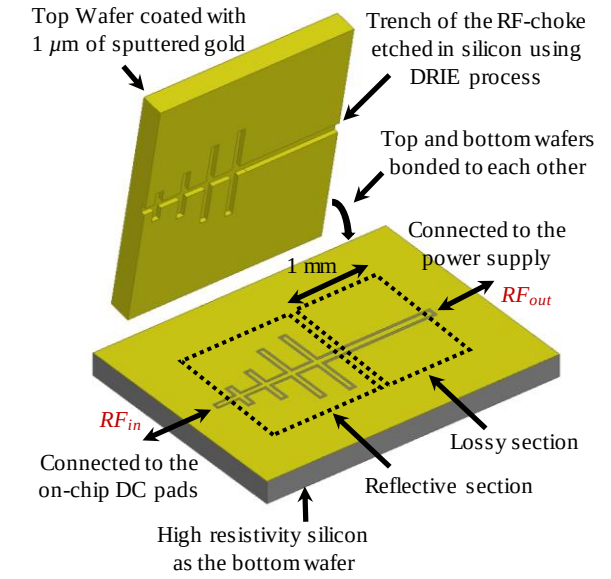
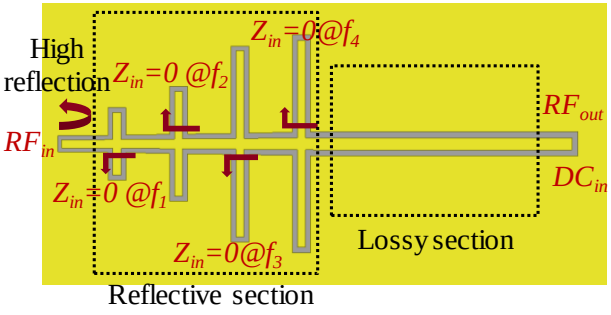


Fig. 5. Full-wave simulated response of the proposed RF-choke showing broadband (40-330 GHz) performance of the RF-choke.



(a)



(b)

Fig. 4. Configuration of the wideband RF-choke incorporated in the chip packaging block, (a) 3-D view, and (b) top view showing implementation of the RF-choke with a series of open-ended stubs with optimized dimension each tuned at different center frequencies.

C. Broadband RF-Choke Design

A wideband planar RF-choke compatible with silicon micromachining is designed to block the loading effects of the DC bias network from the RF circuit components and avoid the possible instability of the integrated active devices. The proposed RF-choke is composed of two sections, namely, the reflective and the lossy sections. The reflective part of the RF-choke is intended to block the RF signal by providing a high RF impedance. This is achieved by designing a series of open-ended stubs [as shown in Fig. 4(b)] implemented on a CBCPW line which is patterned on a high-resistivity silicon wafer. Fig. 4 illustrates the configuration of the designed RF-choke. The dimensions of the stubs are optimized using a full-wave solver (HFSS) to achieve a high RF reflection over a wide bandwidth. The lossy section of the RF choke is implemented by an extended section of the CPW line (~1 mm) and is intended for a better RF isolation by further suppressing any leaked RF signal that has passed through the reflective

section. Fig. 5 shows the full-wave simulation of the proposed RF-choke showing an isolation of better than 10 dB with a reflection loss of less than 2 dB over the entire desired frequency range (40-330 GHz).

D. Packaging Sensitivity Analysis

Accurate alignment of the MMIC and the on-package RF and DC lines is very challenging and without a reliable method it can be the major error source in measurements. In our approach, the provision of the mounting block on the bottom side of the package restricts the possible sources of misalignment (rotation and lateral displacement) significantly. This window limits the probe positioning to a maximum of 10 µm. In addition to the misalignment errors, the packaging is also prone to non-ideal contact between the RF pads on the package and the on-chip RF ports. The gaps (up to 2 µm) can degrade the RF performance of the transition to some extent and its effect on the RF performance of the packaging scheme needs to be studied. In order to validate the performance of the package with respect to the aforementioned sources of errors (dx, dy) [Fig. 6 (a)], full-wave simulation study of the package is carried out and the results are presented here. Figure 6(b) shows the sensitivity analysis of the package for different gap values between the packaged MMIC and the on-wafer lines. As shown, for gap values as large as 2.5 µm, the gap has minimal effect (less than 0.4 dB of loss) on the performance of the packaged MMIC. Figure 6(c) also shows the sensitivity analysis of the package for possible misalignments between the on-chip pads and the on-wafer package lines, where misalignments as large 10 µm and 5 µm in y and x directions, respectively, only increase the insertion loss of the package by 1 dB or less.

III. Microfabrication Process

The fabrication of the proposed integration block involves a high level of fabrication accuracy which is difficult and costly to achieve using the conventional milling techniques. On the other hand, silicon micromachining offers accurate fabrication of complex structures and has the advantage of having low mass and less cost. This makes silicon micromachining a suitable option for fabrication of the packaging block along with the rest of the radar system on a single silicon substrate.

The microfabrication of the peg and socket chip integration block is accomplished using two silicon wafers, namely, top and bottom wafers. The bottom wafer has a thickness of 1000 µm and contains the trench for the CBCPW line and perhaps waveguides and antennas as well as a thru-wafer opening for mounting the chip. The trenches and the opening in the bottom-wafer are realized by etching the silicon wafer by using the DRIE Bosch process from the top and bottom sides, respectively. The bottom wafer is next fully metalized by 2 µm of sputtered gold. The top wafer contains the patterns of the CBCPW traces and the transition to the RF pads which are patterned over a 1 µm thick layer of SiO_2–Si_3N_4–SiO_2 membrane deposited over it. The topside of the top wafer over the CBCPW line and the opening for the chip mounting is also etched down to the membrane layer using the DRIE silicon etch process. The bottom and top wafers are then aligned and bonded using gold-to-gold thermocompression bonding. A close-up of the fabricated sample is shown in Fig. 7.

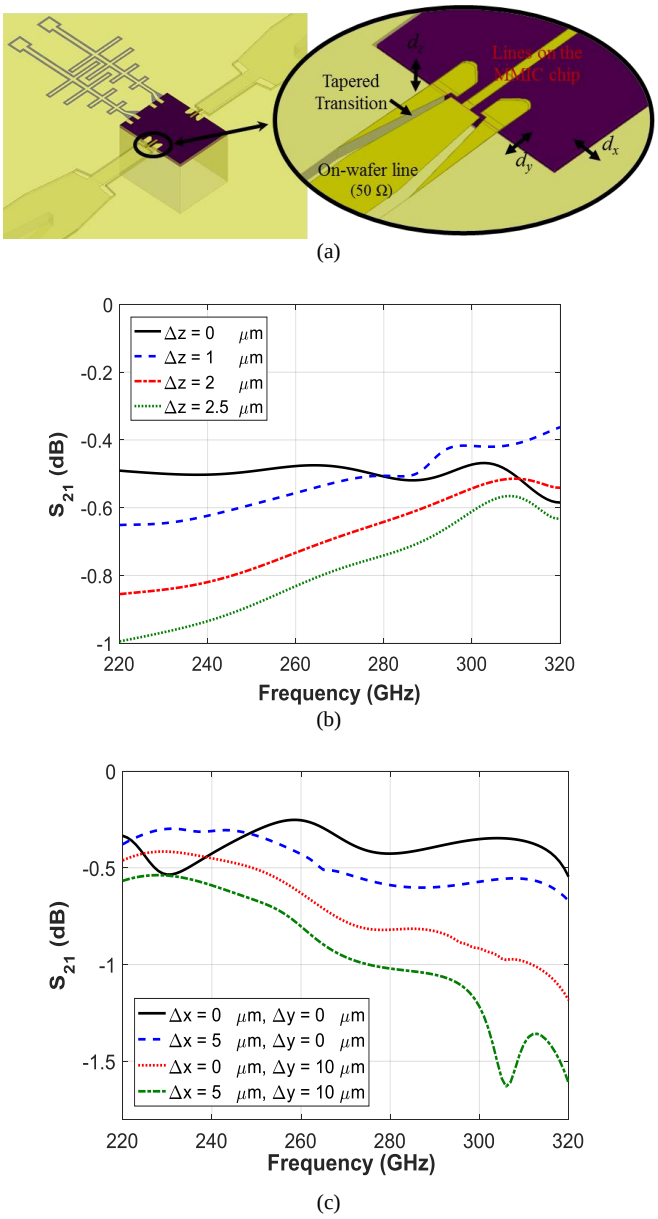


Figure 6. (a) Schematic diagram of the package with possible misalignments (d_x , d_y) and gaps (d_z) between the on-wafer pads and those of the MMIC, (b) full-wave simulation of the effect of the gaps on the RF performance of the package, and (c) sensitivity analysis of the RF performance of the package with respect to misalignments over the entire WR-3 band.

IV. Measurement Process

A. DC Measurements

In order to verify the biasing condition of the packaging block, DC measurements of a packaged chip is performed using an Agilent U2723A source measurement unit. In this measurement a single stage amplifier chip is integrated into the packaging block [Fig. 8(a)] and the I-V characteristics of the transistor is extracted, as presented in Fig. 8(b), where

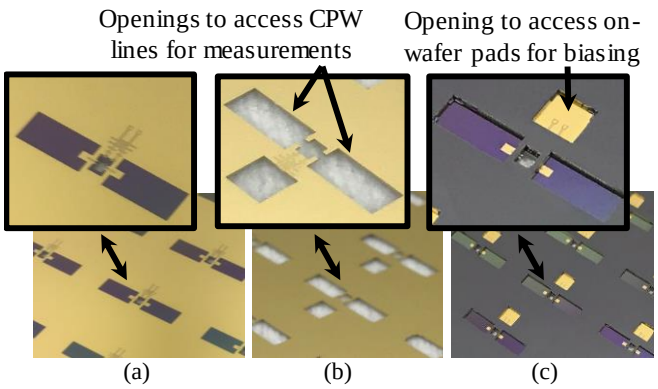


Fig. 7. Close-up of the micromachined packaging block, (a) bottom wafer containing the RF and DC lines patterned on a high resistivity silicon wafer, (b) openings for mounting the MMIC as well as access to the CPW lines etched on a gold coated wafer as the top wafer, and (c) top and bottom wafers bonded to each other using thermocompression bonding process.

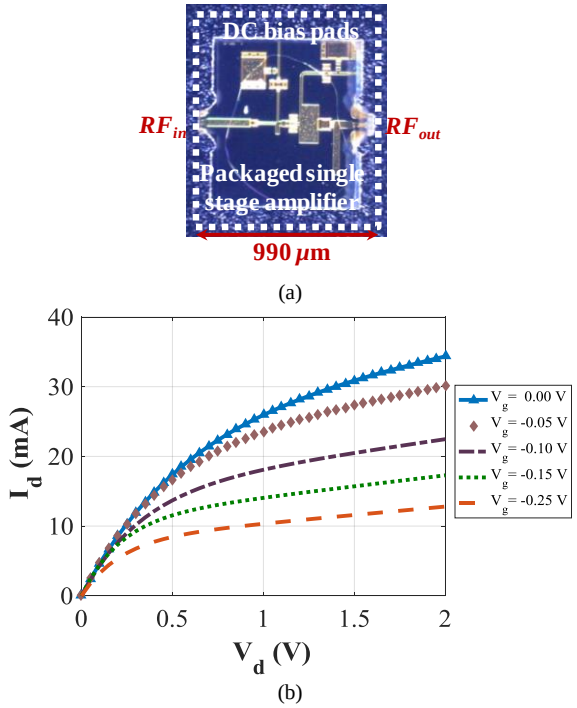


Fig. 8. (a) Micrograph of a packaged amplifier operating at 210 GHz, and (b) Measured I-V characteristics of the packaged device.

successful biasing of the device is shown. Additionally, the DC Ohmic resistance between the on-wafer DC lines and the on-chip DC pads is measured, after de-embedding the resistance of the probes, and a resistance of 2.7 Ω per contact is reported.

B. RF Measurements of the Broadband RF-Choke

To characterize the performance of the RF-choke over a wide frequency range, a setup consisting of an Agilent N5245 4-port network analyzer along with OML MMW frequency extending modules with the appropriate Ground-Signal-Ground (GSG) probes is used to enable two-port measurements from DC to 330 GHz with the exclusion of 40-70 GHz and 120-200 GHz frequency bands. Fig. 9 illustrates the measured *S*-parameters of the RF-choke where a reflection loss of better than 2 dB with an RF isolation of better than 15 dB is achieved over the entire frequency range (30-330 GHz).

C. RF Measurements of Thru-line Chip

To characterize the performance of the packaging block, a thru-line chip is integrated into the block [Fig. 10(a)] and the *S*-parameters of the chip are measured using the measurement setup with WR-3 frequency extender units. The packaged thru-line chip exhibits reflection losses of better than 10 dB with an insertion loss of better than 1 dB per transition over the entire desired frequency range, as shown in Fig. 10(b).

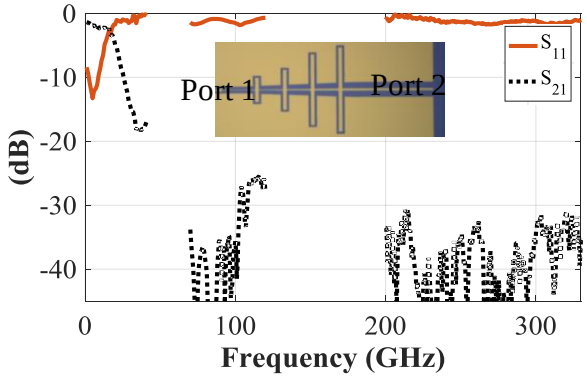


Fig.9. Measured *S*-parameters of the RF-choke over multiple frequency bands verifying the wideband (30-330 GHz) response of the RF-choke with reflection loss of better than 2 dB.

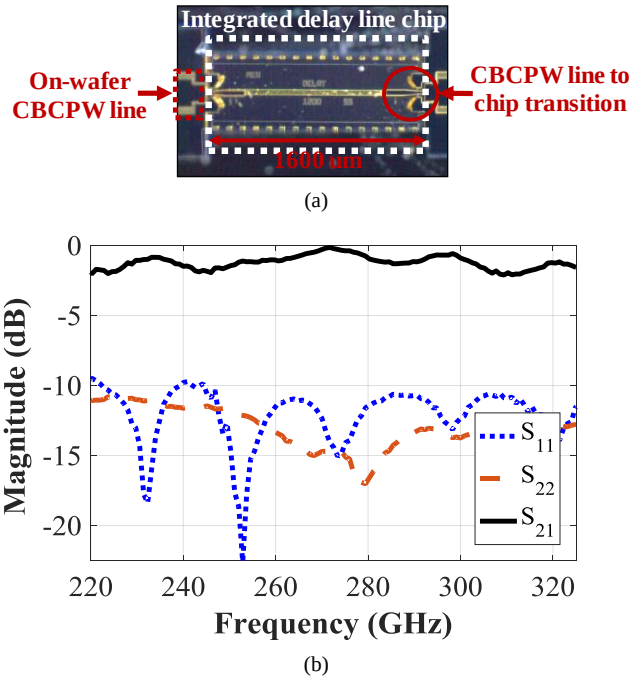


Fig. 10. (a) Close-up of a thru-line chip packaged into the block, and (b) Measured back-to-back *S*-parameters of the packaged chip showing an insertion loss of better than 2 dB over the entire Y-band.

D. RF Measurements of Active MMIC

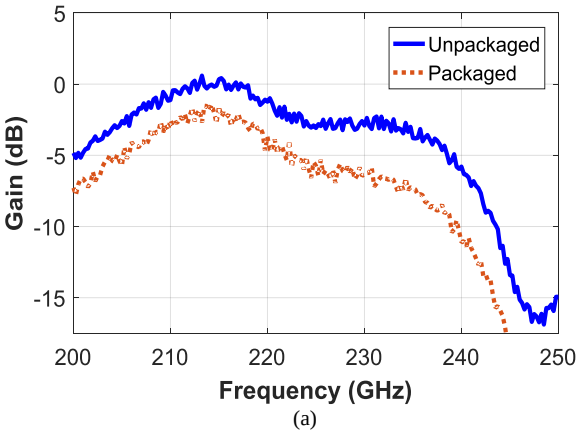
To characterize the performance of the packaging block for integration of active MMIC, a single stage amplifier (BAE Systems) with a poor gain (loss) operating at 210-220 GHz frequency range is integrated into the block [Fig. 8(a)] and the S-parameters of the packaged amplifier are measured using the measurement setup with WR-3 frequency extender units. Figure 11 illustrates the measured *S*-parameters of the MMIC before and after packaging. As shown, the package introduces minimal effect (less than 1.2 dB of loss per transition) on the performance of the packaged MMIC.

CONCLUSION

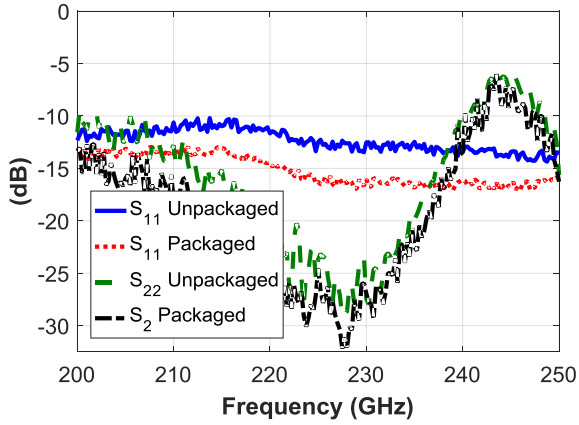
A novel packaging method for integration of active MMIC for MMW to THz band applications has been proposed and demonstrated. The packaging block and the associated biasing circuitry are designed to be compatible with silicon micromachining technology for high level of fabrication, low cost, and small form-factor. The high accuracy and repeatability of the fabrication technology also enables scalability of the packaging method up to THz frequencies, where is expected to find a wide range of applications in this area. In the next quarter the active MMICs are going to be integrated with the rest of the radar modules to complete the radar prototype.

ACKNOWLEDGMENT

This work was supported by the U.S. Army Research Laboratory under Contract W911NF and prepared through collaborative participation in the Microelectronics Center of Micro Autonomous Systems and Technology (MAST) Collaborative Technology Alliance (CTA).



(a)



(b)

Figure 11. Measured *S*-parameters of a packaged single stage amplifier with poor gain (loss), before and after packaging. (a) Measured gain (S_{21}) of the amplifier showing minimal effect of the package (<1.2 dB per transition) on the performance of the MMIC, and (b) Measured reflections of the MMIC verifying minimal effects of the packaging technique.

REFERENCES

[1] R. Lai *et al*., "Sub 50 nm InP HEMT device with Fmax greater than 1 THz", *Proc. IEEE IEDM Conf. Dig.*, pp. 609-611, 2007-Dec.

[2] W. Deal, *et al*., "THz Monolithic Integrated Circuits Using InP High Electron Mobility Transistors," *IEEE Transactions on Terahertz Science and Technology*, vol. 1, no. 1, pp. 25-32, Sept. 2011.

[3] K. Sarabandi, M. Vahidpour, and M. Moallem, "Compact beam scanning 240 GHz radar for navigation and collision avoidance," *Proc. SPIE*, vol. 8031, pp. 803113-1-803113-8, Apr. 2011.

[4] I. Kallfass *et al*., "All active MMIC-based wireless communication at 220 GHz," *IEEE Trans. THz Sci. Tech.*, vol. 1, no. 2, pp. 477–487, Nov. 2011.

[5] A. Jam, M. Moallem, J. East and K. Sarabandi, "A Non-Contact Waveguide Probe for On-Wafer S -Parameter Measurements for Submillimeter-Wave to Terahertz Band," *IEEE Transactions on Terahertz Science and Technology*, vol. 4, no. 4, pp. 515-522, July 2014.

[6] A. Jam and K. Sarabandi, "A horizontally polarized beam-steerable antenna for submillimeter-wave polarimetric imaging and collision avoidance radars," *2016 IEEE Int. Symp. on Antennas Propagat*, Fajardo, 2016, pp. 789-790.

[7] A. Jam, M. Vahidpour, J. East and K. Sarabandi, "A sub-millimeterwave micromachined frequency beam-steering antenna array," *2014 IEEE MTT-S International Microwave Symposium (IMS2014)*, Tampa, FL, 2014, pp. 1-3.

[8] L. Samoska, "An overview of solid-state integrated circuit amplifiers in the submillimeter-wave and THz regime," *IEEE Trans. THz Sci. Tech*, vol. 1, no. 1, pp. 9–24, Sept. 2011.

[9] K. Leong *et al*., "A 340 GHz integrated CB-CPW-to-waveguide transition for sub millimeter-wave MMIC packaging," *IEEE Microw. Wireless Compon. Lett.*, vol. 19, no. 6, pp. 413–415, Jun. 2009.

[10] M. Moallem *et al*., "A Broadband, Micromachined Rectangular Waveguide to Cavity-Backed Coplanar Waveguide Transition Using Impedance-Taper Technique," *IEEE Trans. THz Sci. Tech*., vol.4, no.1, pp.49,55, Jan. 2014.



Armin Jam (S'09) received the B.Sc. degree in electrical engineering from Shahid Beheshti University, Tehran, Iran, in 2010, and the M.Sc. and Ph.D. degrees in electrical engineering from The University of Michigan at Ann Arbor in 2013 and 2017, respectively.

He has been a research assistant with the Radiation Laboratory, University of Michigan, since 2012 where he has been developing millimeter-wave (MMW) and terahertz devices. His research interests include MMW and sub-MMW antennas and radars; high-frequency instrumentation and metrology; as well as micromachining technology of high-frequency devices.

Dr. Jam was the recipient of the Best Student Paper Award in the 2016 IEEE International Symposium on Antennas and Propagation (AP-S) and an Honorable Mention in the 2017 URSI General Assembly and Scientific Symposium. He has received the University of Michigan Richard F. and Eleanor A. Towner prize for distinguished academic achievement in 2014 and was a winner of Engineering Graduate Symposium Technical Award from the College of Engineering in 2013 and 2016. He was also a recipient of the Silver Medal in the National Mathematics Olympiad in 2005. He is a member of Phi Kappa Phi, Tau Beta Pi, and the National Elites Foundation, and is listed in Who's Who in America.



Jack East (LM'13) received the B.Sc. and Ph.D. degrees from The University of Michigan.

He is currently with the Electrical Engineering and Computer Science Department, The University of Michigan, where he conducts research in the areas of high-speed microwave device design and fabrication, circuit modeling and characterization and THz devices.



Kamal Sarabandi (S'87-M'90-SM'92-F'00) received the B.S. degree in electrical engineering from the Sharif University of Technology, Tehran, Iran, in 1980, the M.S. degree in electrical engineering in 1986, and the M.S. degree in mathematics and the Ph.D. degree in electrical engineering from The University of Michigan at Ann Arbor in 1989.

He is currently the Director of the Radiation Laboratory and the Rufus S. Teesdale endowed Professor of Engineering in the Department of Electrical Engineering and Computer Science, The University of Michigan at Ann Arbor. His research areas of interest include microwave and millimeter-wave radar remote sensing, Meta-materials, electromagnetic wave propagation, and antenna miniaturization. He possesses 25 years of experience with wave propagation in random media, communication channel modeling, microwave sensors, and radar systems and leads a large research group including two research scientists and 16 Ph.D. students. He has graduated 46 Ph.D. and supervised numerous post-doctoral students. He has served as the Principal Investigator on many projects sponsored by the National Aeronautics and Space Administration (NASA), Jet Propulsion Laboratory (JPL), Army Research Office (ARO), Office of Naval Research (ONR), Army Research Laboratory (ARL), National Science Foundation (NSF), Defense Advanced Research Projects Agency (DARPA), and a large number of industries. Currently he is leading the Center for Microelectronics and Sensors funded in 2008 by the Army Research Laboratory under the Micro-Autonomous Systems and Technology (MAST) Collaborative Technology Alliance (CTA) program. He is also leading a newly established center in Microwave Sensor Technology funded by King Abdulaziz City for Science and Technology (KACST).

He has published many book chapters and more than 250 papers in refereed journals on miniaturized and on-chip antennas, meta-materials, electromagnetic scattering, wireless channel modeling, random media modeling, microwave measurement techniques, radar calibration, inverse scattering problems, and microwave sensors. He has also had more than 580 papers and invited presentations in many national and international conferences and symposia on similar subjects.

Dr. Sarabandi served as a member of NASA Advisory Council appointed by the NASA Administrator for two consecutive terms from 2006-2010. Currently he is serving as the President of the IEEE Geoscience and Remote Sensing Society (GRSS). He was a member of the Editorial Board of the PROCEEDINGS of the IEEE and an associate editor for the IEEE TRANSACTIONS ON ANTENNAS AND PROPAGATION and the IEEE Sensors Journal. He is a member of Commissions F and B of URSI and is serving as the vice Chair of the USNC URSI Commission F. Dr. Sarabandi was the recipient of the Henry Russel Award from the Regent of The University of Michigan. In 1999 he received a GAAC Distinguished Lecturer Award from the German Federal Ministry for Education, Science, and Technology. He was also a recipient of the 1996 EECS Department Teaching Excellence Award and a 2004 College of Engineering Research Excellence Award. In 2005 he received the IEEE GRSS Distinguished Achievement Award and the University of Michigan Faculty Recognition Award. He also received the best paper Award at the 2006 Army Science Conference. In 2008 he was awarded a Humboldt Research Award from The Alexander von Humboldt Foundation of Germany and received the best paper award at the IEEE Geoscience and Remote Sensing Symposium. He was also awarded the 2010 Distinguished Faculty Achievement Award from the University of Michigan. The IEEE Board of Directors announced him as the recipient of the 2011 IEEE Judith A. Resnik Award. Professor Sarabandi was recognized by the IEEE GRSS with its 2013 Education Award. In 2016 he received the Stephen S. Attwood Award from the College of Engineering at the University of Michigan and a Distinguished Alumni Award from Sharif University of Technology. In the past several years, joint papers presented by his students at a number of international symposia (IEEE APS'95,'97,'00,'01,'03,'05,'06,'07,'16; IEEE IGARSS'99,'02,'07;'11,'14, IEEE IMS'01, USNC URSI'04,'05,'06,'10,'11, AMTA '06, URSI GA '08,'14; Eastern Snow Conference '16) have received best paper awards.

Through-the-wall Radar Imaging Using an Ad hoc Ensemble of MAST Robotic Platforms

Kamal Sarabandi, Behzad Yektakhah, and Xiuzhang Cai

***Abstract*—Through-the-wall radar imaging systems are utilized for mapping buildings' interiors and detecting static and moving objects hidden behind the walls. The current techniques rely on large linear antenna arrays with directional radiation pattern for obtaining high cross-range resolution, and their imaging capability is limited by the low array processing gain and field of view. In this project, a new technique for through-the-wall radar imaging in which the linear array is replaced by a dense 2-D synthetic array formed by an ad hoc network of moving transceivers mounted on MAST robots is introduced. Applying this method enables imaging of building interiors from outside or inside. Utilizing omnidirectional antennas, the method can provide 360° image of the surrounding environment in a short time. In this report, first, the proposed imaging technique is briefly reviewed and then the tasks that have been done during December 2015 to December 2016 are described. The tasks are focused on fabrication and test of a bi-static FMCW radar system to realize the proposed method for the through-the-wall radar imaging. The system includes a fast chirp signal generator (FMCW transmitter), FMCW receiver, wireless synchronization circuits, wideband circularly polarized omnidirectional antennas, and digital circuits (including A/D and data acquisition circuits). At the end of this report, a 360° image formed by the fabricated system is presented.**

***Index Terms*—Bi-static FMCW, circularly polarized antenna, MAST, synthetic aperture radar, through-the-wall imaging.**

I. INTRODUCTION

ONE of the important capability requirements mentioned in the SPP sub-thrusts is the intelligence, surveillance, and reconnaissance (ISR). Another important required capability for the autonomous MAST platforms is navigation, mapping, and path planning. Sensors that enable both requirements are of high value and importance for autonomous operation of small robotic platforms. Selection of any sensor for small robots is constrained by SWaP and environmental conditions. MAST missions may take place in lit, dark, or optically obscured environments and therefore it is highly desired that the sensor be independent of the ambient condition such as external light source. Mapping the building layout from outside and finding possibly obscured ingress and egress points can significantly enhance the operation of an autonomous robotic ensemble exploring a building. In addition, threat assessment in terms of determining whether the building is occupied or not is of paramount importance for their missions. Similar functions are also important once the robots are inside the building and are exploring other rooms from a given location inside a building or searching for hidden rooms or objects behind the walls.

The authors are with the Radiation Laboratory, Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109-2122 USA (email: byekta@umich.edu; saraband@eecs.umich.edu).

Through-the-wall radar imaging using EM waves are intended for mapping buildings' interiors and detecting static and moving objects hidden behind walls and other optically opaque obstacles. This capability has found many applications in rescue, military, and law enforcement operations.

Radars considered for through-the-wall imaging applications make use of ultra wideband signal waveforms that typically fall in the frequency range 0.5−2 GHz to provide the required range resolution. This frequency band is low enough to allow for signal penetration through typical walls and high enough to allow for physically small efficient antennas for the system. To achieve high cross-range resolution in the final image, a large aperture must be formed. To achieve this, two different approaches may be pursued: 1) using a large array of transmitting and receiving antennas (real aperture) and 2) using a moving transmitter and receiver pair to realize a large synthetic aperture. For achieving fine spatial resolution, the use of real aperture antennas is impractical. Also, the performance of the standard drive by synthetic aperture systems are limited due to the limited field of view of such systems. These systems for mapping the buildings from outside can usually provide image of the first floor of the buildings and their imaging capability decreases as the complexity of the buildings' interiors increases. With directional antennas, conventional imaging systems provide only image of the area in front of the antennas and therefore they require multiple scans from different directions, if possible, to provide a 360° view of the buildings' interiors.

In this project, a new approach for through-the-wall radar imaging from within buildings is proposed. The proposed method allows for improved mobility within confined space without increasing complexity and provides a complete view of the building interiors in a short period of time. The main idea is to replace the array of directive antennas moved on a linear path with a limited number of small moving robotic transceivers with omnidirectional antennas as illustrated in Fig. 1. In this method, one robot acts as transmitter and the rest act as receivers. Receiving robots move inside the room and sample the reflected signal from the building's interiors and hidden objects at different positions (sampling points). Utilizing

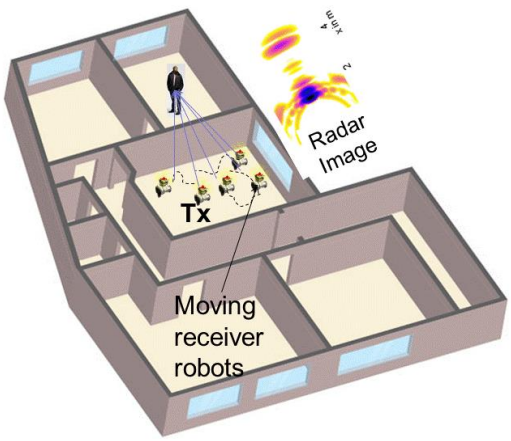


Fig. 1. Transmitter (Tx) and receiver (Rx) arrangement in the proposed method.

multiple receiving robots enables obtaining a large number of samples in a short time. Using omnidirectional transmitting and receiving antennas and applying an appropriate beam forming technique to the received signals, a large and dense 2-D array can be synthesized to form a 360° view of the imaging area with a very high cross-range resolution and high processing gain.

In this report, the proposed imaging technique is briefly reviewed in Section II. A bi-static FMCW transmitter and receiver which are designed and fabricated to realize the method is also introduces in Section II. The tasks that have been done during December 2015 to December 2016 are described in detail in the next sections. The tasks are focused on fabrication and test of a bi-static FMCW radar system (excluding the RF transmitter and receiver) to realize the proposed method for the through-the-wall radar imaging. The system includes wireless synchronization circuits, wideband circularly polarized omnidirectional antennas, and digital circuits (including A/D and data acquisition circuits) which are described in Sections III−V. In Section VI, preliminary imaging results are presented. Final imaging results using the fabricated antennas are presented in Section VII. Concluding remarks are given in Section VIII.

II. Review of the Tasks Before December 2015

Fig. 2(a) shows the geometry of a simple through-the-wall radar imaging problem and configuration of transmitter and receivers according to the proposed method. According to Fig. 2(a), one stationary transmitter and one or more moving receivers are placed within a room with dielectric walls. Utilizing omnidirectional antennas, the transmitter sends a signal in all directions. Receivers move inside the room and sample the reflected signal from the building's walls and hidden objects as well as the direct signal from the transmitter at different positions. Each sampling spot corresponds to an antenna element of a large synthetic array. To determine the relative locations of the transmitter and moving receivers, a flying robot equipped with a high-resolution camera and perched at a location that can view the scene is used. This is illustrated in Fig. 2(b).

To form the image, the standard back propagation technique

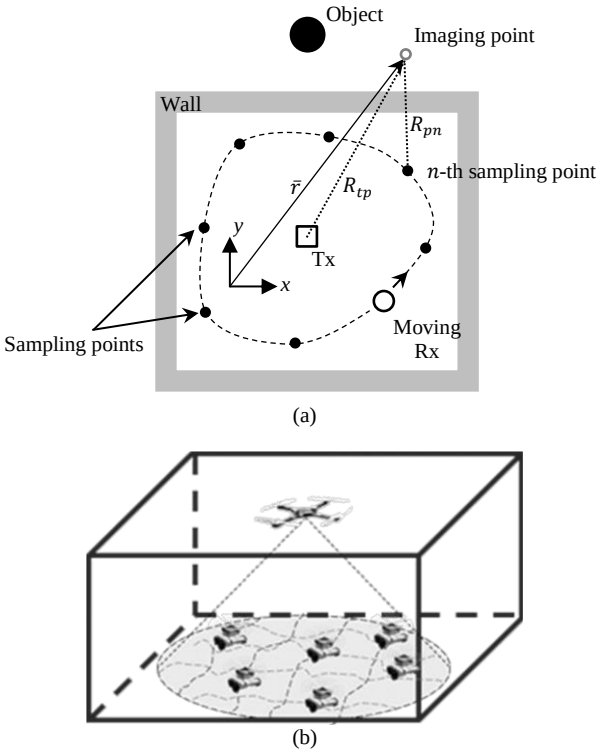


Fig. 2. (a) Geometry of a simple through-the-wall radar imaging problem. (b) Flying robot for detection of the positions of Tx and Rxs.

is applied to the samples of the received signal. Considering N_f frequency samples and N_s sampling points, the image at position $\bar{r}$ ($S(\bar{r})$) is calculated by:

$$S(\bar{r}) = \frac{1}{N_f N_s} \sum_{m=1}^{N_f} \sum_{n=1}^{N_s} a_n S_{mn}^r R_{tp} R_{pn} e^{jk_m(R_{tp}+R_{pn})} \quad (1)$$

where, S_{mn}^r is the received signal at m-th frequency sample and n-th sampling point, R_{tp} is the distance between the transmitter and the imaging point at position $\bar{r}$, R_{pn} is the distance between the imaging point at position $\bar{r}$ and n-th sampling point (Fig. 2), and k_m is the propagation constant at m-th frequency. a_n is the array weight factor for beam forming [1].

To form a 360° image in a short period of time, omnidirectional antennas are utilized in the proposed method. Using omnidirectional antennas arises some problems that must be addressed in the method. The problems are: A) ambiguity in detection of objects' positions and B) receiving multiple interactions of rays with the surrounding environment.

A. *Ambiguity in Detection of Objects' Positions*

Since omnidirectional antennas cannot discriminate between the signals received from different directions, they can produce ambiguity in detection of the real positions of the objects. It is shown in Fig. 3(a) where a linear array of sampling points is used to form the image of two point objects at the top. It can be observed that in addition to the images of the objects, two unwanted images are formed at the bottom. This is due to the fact that the antennas cannot discriminate between the signals received from the top and bottom and also all the sampling

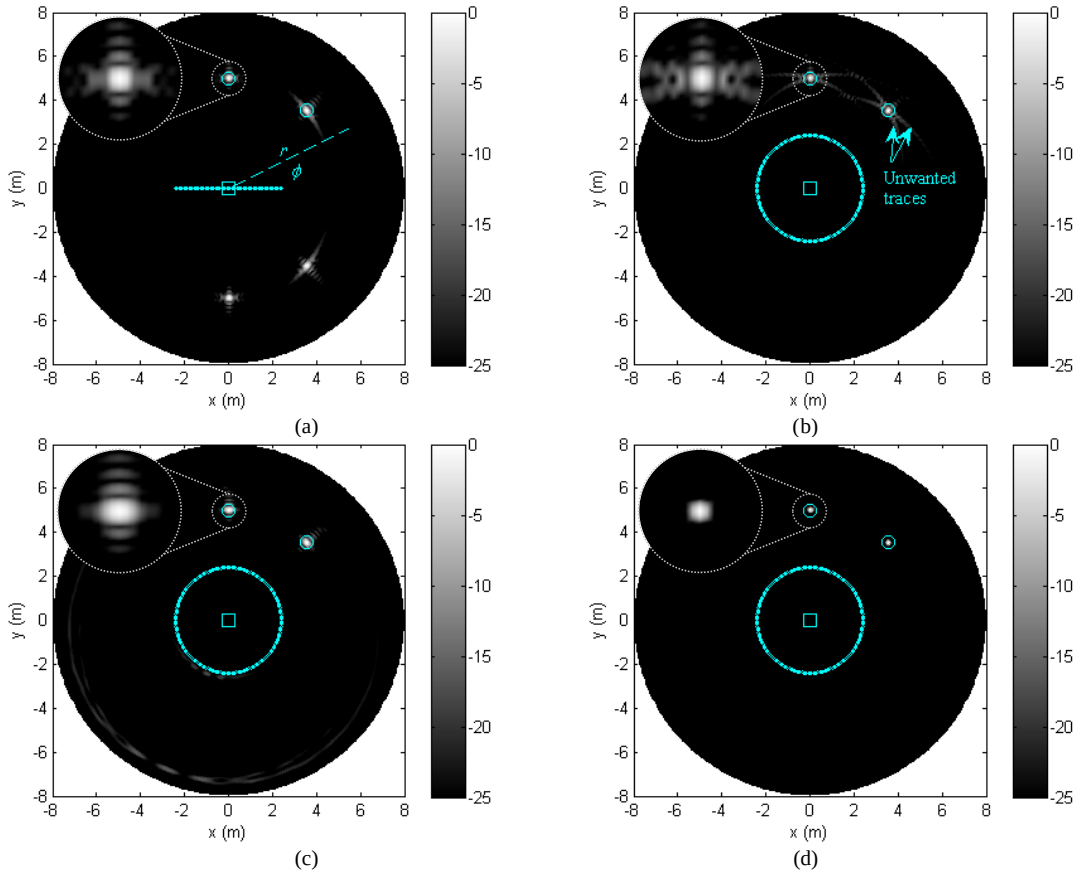


Fig. 3. Normalized image calculated for two point targets located at (r, ϕ) = (5 m, 90°) and (5 m, 45°). (a) Image formed by a linear array of sampling points (using uniform array coefficients). (b) Image formed by an array of sampling points on a circle (using uniform array coefficients). (c) The same image as (b) obtained by cosine-squared coefficients. (d) Multiplication of (b) and (c).

points on the line mirror the real images of the objects at the same position at the bottom. As shown in Fig. 3(b), this ambiguity is removed by using an array of sampling points on concentric circles. In this configuration, since each array element mirrors the image of each object at different position, the final mirrored image points are not coherently added. With this arrangement for sampling points, a uniform cross range resolution can be also obtained at all directions. Images in Fig. 3(a) and (b) are obtained by uniform array weight factors (a_n = 1).

To remove the unwanted traces that are still present in Fig. 3(b), non-uniform array weight factors can be used. These traces are due to the side-lobes in the uniform array. To reduce the sidelobe level, non-uniform array coefficients can be used at the expense of wider beamwidth and lower cross-range resolution. To reduce the sidelobes, cosine-squared array coefficients are used in the proposed method. Fig. 3(c) shows the same image as Fig. 3(b) while the image is calculated by (1) with cosine-squared coefficients. In this case the unwanted traces have been removed while the cross-range resolution has been decreased in all directions. To maintain the fine cross-range resolution of the uniform array and the low sidelobe level of the cosine-squared array in the final image, the two images in Fig. 3(b) and Fig. 3(c) can be multiplied. The result is shown in Fig. 3(d) where it is shown that unwanted traces are reduced while the cross-range resolution is approximately the same as the Fig. 3(b).

B. Receiving Multiple Interactions of Rays with the Surrounding Environment

Images of the walls and objects are formed by the first reflections of the transmitted wave. The higher order interactions produce spurious images in the image and increases the background noise. The significant high order interactions are double reflection from the corners and double reflection from the parallel walls. Omnidirectional antennas can receive these higher order interactions and thus the spurious images appear in the final image. Fig. 4(a) shows the image of a room with four walls and an object behind the walls while vertical polarization is used for both transmitter and receiver. The effect of different interactions is indicated in the figure. Omnidirectional antennas receive the direct signal from transmitter to receivers. This relatively high power signal can saturate the receivers. The peak at the center of the room in Fig. 4(a) is due to the direct signal. To reduce the effect of the direct signal and even number of reflections orthogonal circular polarizations are utilized for transmitter and receivers. In this case, direct signal and even number of reflections have the same polarization as the transmitted wave and therefore they are attenuated by polarization isolation of the receiving antenna. Fig. 4(b) shows the image obtained by LHCP for the transmitter and RHCP for receivers. It can be observed that effects of even number of reflections and direct signal are significantly reduced.

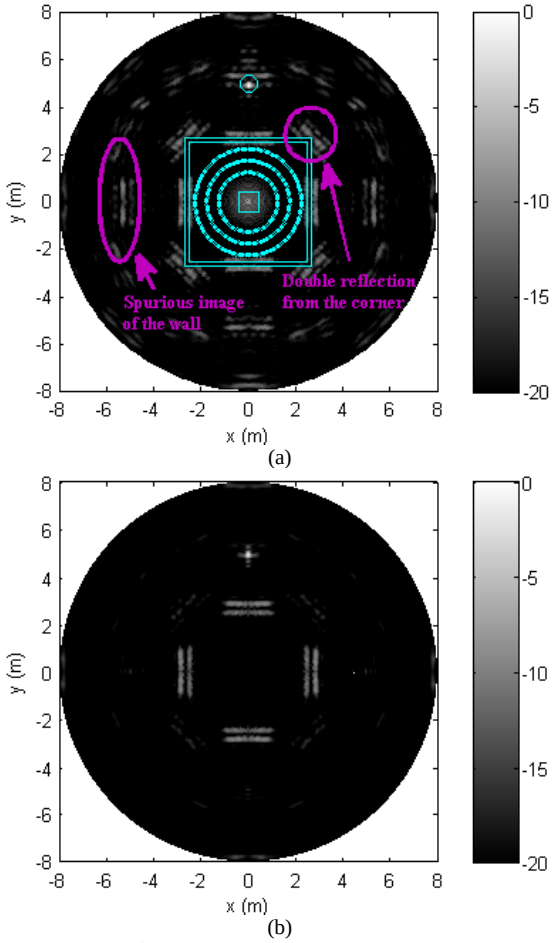


Fig. 4. (a) Geometry of the problem where the wall boundary is shown by solid square, Tx position is shown by a small square, target position is shown by a small circle, and sampling points are shown by dots. Also shown is the normalized image (dB) for Tx and Rx antennas with vertical polarization. (b) Normalized image (dB) for orthogonal circular polarizations (LHCP for transmitter and RHCP for receiver).

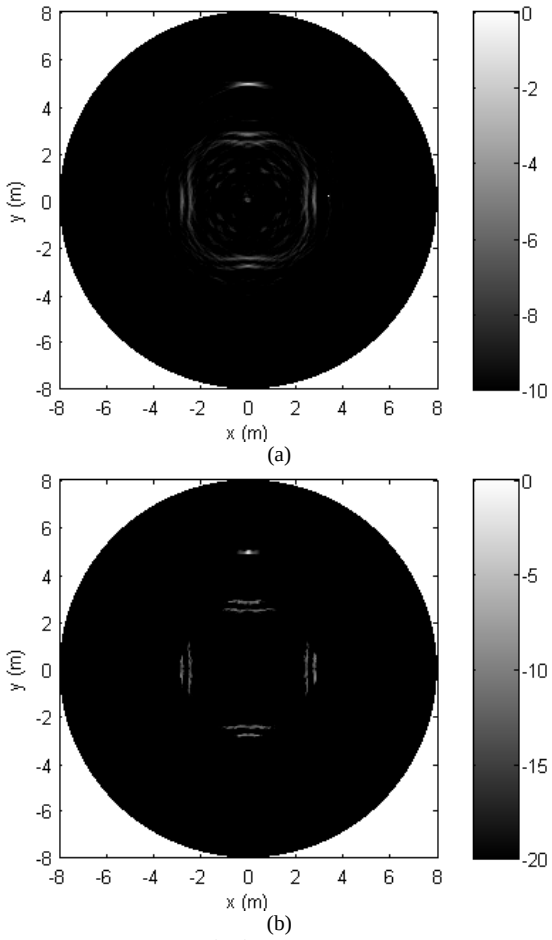


Fig. 5. (a) Detected locations of reflections in Fig. 4(b). (b) The same image as Fig. 4(b) except that it is calculated only at the detected locations of reflections in (a).

To further improve the image, a method based on generalized pencil of function (GPOF) is proposed. With this method, the locations of reflections which are wall and objects boundaries are detected and then the imaging is only performed at the detected locations. Applying the proposed method decreases the background noise and enhances the range resolution of the image. Fig. 5(a) shows the detected locations of reflections in Fig. 4(b). Fig. 5(b) is the same as Fig. 4(b) except that it is only calculated at the locations of reflection in Fig. 5(a). It can clearly be observed that the wall boundaries are imaged with a higher range resolution and also the background noise is significantly reduced.

To test the performance of the proposed imaging method in real situations, a bi-static FMCW radar was designed and fabricated.

Fig. 6 shows the final transmitter/receiver which its different components are integrated into a single board (different parts are indicated in the figure). The board is the same for both the transmitter and receiver expect that the mixer part is only soldered for the receivers. The phase noise of the fabricated radar is better than -100 dBc at 18.75 kHz offset over the band 1.2–2 GHz. Total sweep time for the chirp signal (1.2–2 GHz) is 47 μs.

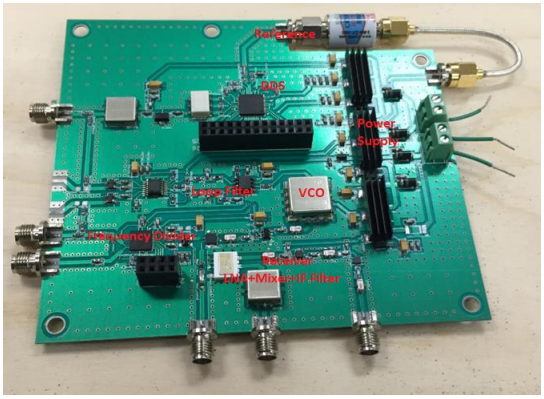
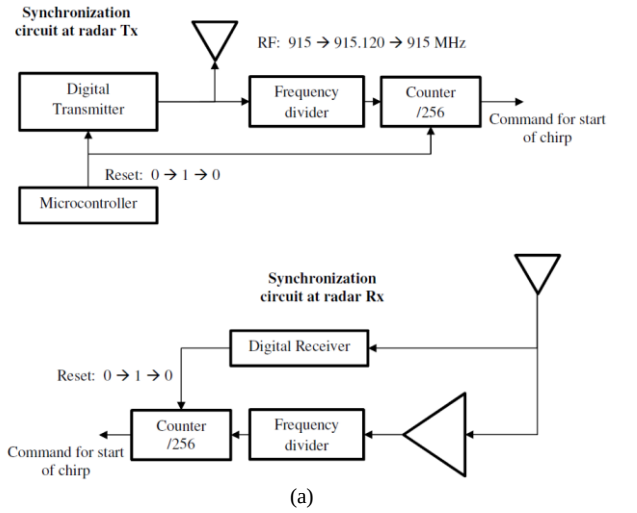


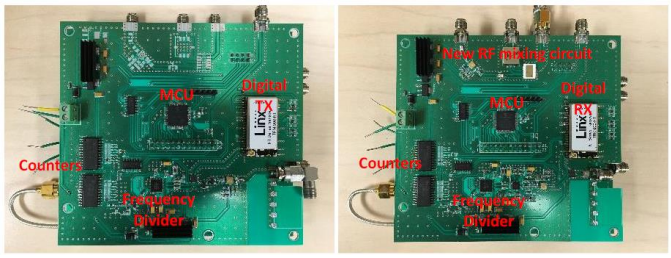
Fig. 6. Fabricated FMCW transmitter and receiver

III. Design, Fabrication, and Test of the Synchronization Circuits

In bi-static FMCW radar, the start time for the chirp signal at both Tx and Rx must be synchronized. Fig. 7(a) shows the block diagram of the proposed synchronization method. The digital transmitter at radar Tx continuously transmits a signal with the frequency of 915 MHz (transmitted data = 0) or 915.080 MHz (transmitted data = 1). This signal is received at Rx. The frequency of the signal is reduced by frequency dividers and counters at both Tx and Rx. The resulting low frequency signal



(a)



(b)

Fig. 7. (a) Block diagram of the proposed synchronization method. (b) Fabricated synchronization circuits for Tx (left) and Rx (right).

is used as a command for the start of the chirp signal. With this configuration, any change in the frequency of the transmitted signal due to the instability of the oscillator in the digital transmitter and change of the input data has the same effect on the generated low frequency signals at both Tx and Rx. After power up, to synchronize the counters at both Tx and Rx, a microcontroller pin at radar Tx which is connected to the reset pin of counters at Tx and also the input of the digital transmitter, changes from 0 to 1 and then from 1 to 0. This transition directly resets the counters at Tx and also resets the counters at Rx via the wireless channel provided by the digital transmitter and receiver. With the proposed method, the maximum time error in synchronization is less than 400 ns. This error can be detected and compensated in IF signal by knowing the distance between the transmitter and receivers. This error is due to the random delay in digital modulator and demodulator. Fig. 7(b) shows the fabricated synchronization circuits for both Tx and Rx. The circuits have been tested and the results show the maximum time error of 400 ns in synchronization of the start of the chirp signal in successive power ups. A microcontroller (MCU) is incorporated into synchronization circuits for sending the reset command for counters and also programming the DDS and frequency divider in radar RF circuits (Fig. 6).

To improve the IF signal, a new RF LNA is used in the RF mixing circuit of the Rx. The new LNA provides more than 30 dB of the gain over the band 1.2–2 GHz. The new RF mixing circuit is included in the synchronization circuit of the Rx.

IV. Wideband Circularly Polarized Omnidirectional Antennas

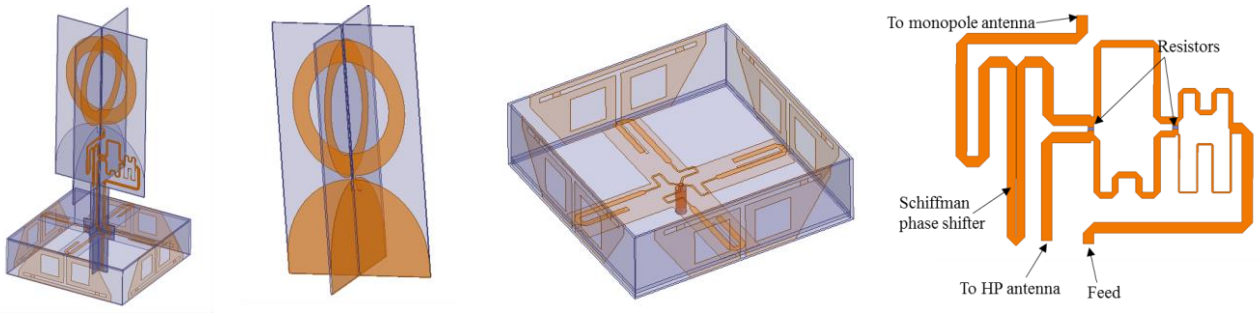
As mentioned in Section II, circularly polarized omnidirectional antennas are required to reduce the effect of even number of reflections and the direct signal from Tx to Rx. The antenna must be wideband to satisfy the system requirement for the range resolution. In this section, two different designs for wideband circular polarized omnidirectional antennas are presented.

A. Wideband Circularly Polarized Omnidirectional Antenna Using Omnidirectional VP and HP Radiating Elements

A broadband omnidirectional circularly polarized antenna design is proposed using omnidirectional VP and HP radiating elements and a proper feed structure. As will be shown an antenna with almost 50% bandwidth for both return loss and 3–dB axial ratio in all azimuth directions has been designed.

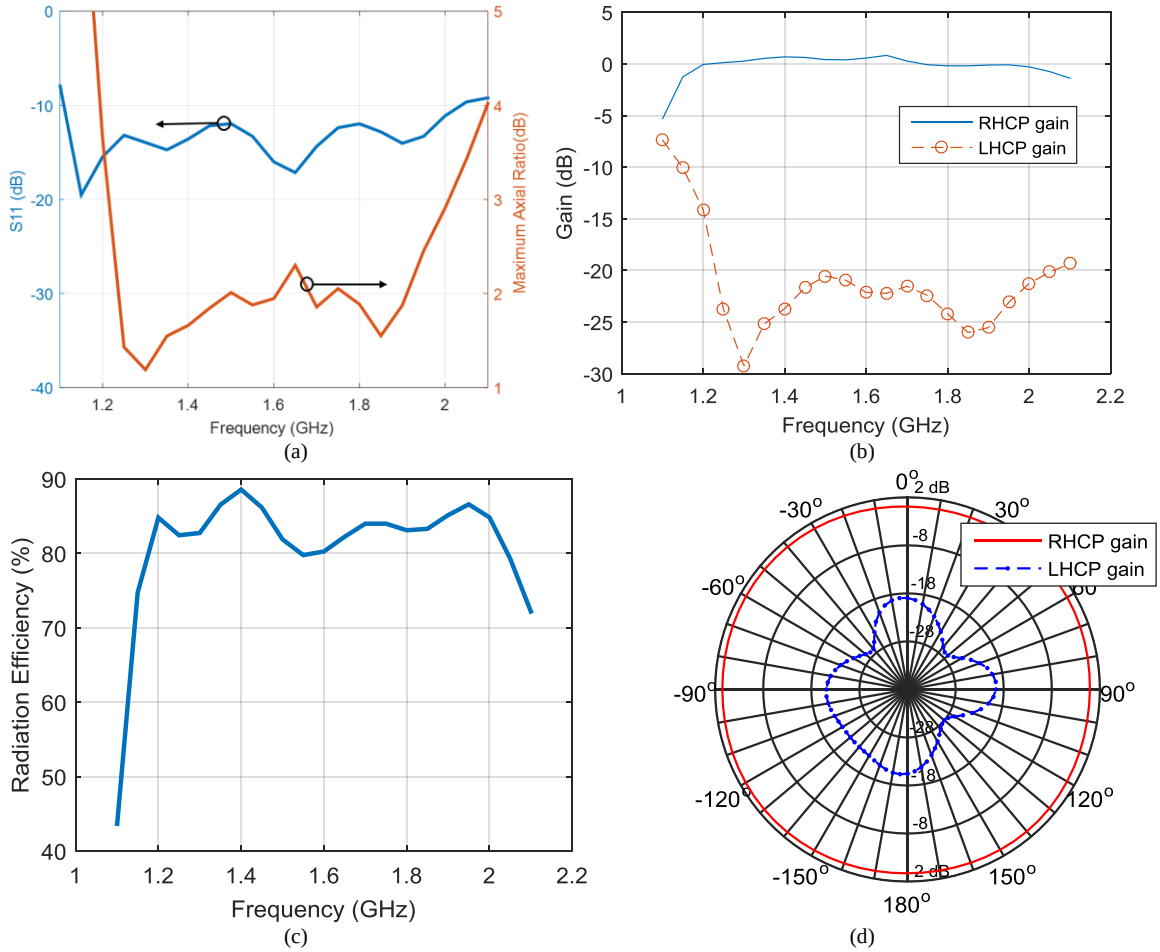
The overall structure of the proposed broadband omnidirectional circularly polarized antenna is shown in Fig. 8(a). It is mainly composed of three parts, a UWB monopole antenna for the VP component, a broadband HP antenna from our previous work, and the feeding network to divide the power with desired ratio and almost 90° phase shift over the entire band. Printed circular monopole antennas have been widely used in UWB applications. The monopole component used in this design is shown in Fig. 8(b). Compared with ordinary printed planar circular monopole antennas, it contains two planar elliptical monopoles intersecting with each other with right angle to obtain better omnidirectional property, and the ground planes are tapered such that the cross polarization can be reduced. The ring shape for the elliptical part is to suppress other radiation modes and to stabilize the impedance as a function of frequency. The compact broadband HP omnidirectional antenna is shown in Fig. 8(c), which is recently developed, also provides more than 50% bandwidth and good omnidirectional property. The feeding network shown in Fig. 8(d) is composed of a two-stage Wilkinson power divider and a Schiffman phase shifter. The main challenge of broadband CP design is to maintain gain and phase difference of two linear polarizations for the entire frequency band and all azimuth directions. The power divider is very carefully designed with unequal impedance and length for two arms to achieve the desired power division ratio as a function of frequency in order to maintain a relatively small gain difference between VP and HP. The Schiffman phase shifter uses a quarter wave coupled line to implement 90°±10° phase shift phase shift in the desired band (1.2−2GHz). There is one extended transmission line between the phase shifter and the feed to the monopole. It is designed to compensate the electric length difference between the path from the feed point of power divider to the HP antenna and that from the feed point to the monopole antenna. In such manner the phase difference between the HP wave and VP wave will be almost constant as 90º for a broadband at all azimuth directions.

The proposed antenna is simulated in HFSS. The return loss result and the maximum axial ratio in all azimuth directions are shown in Fig. 9(a). It can be seen that the 10–dB return loss



(a) (b) (c) (d)

Fig. 8. Structure for (a) the overall CP antenna, (b) the monopole VP antenna component, (c) the HP antenna component, and (d) the feeding network.



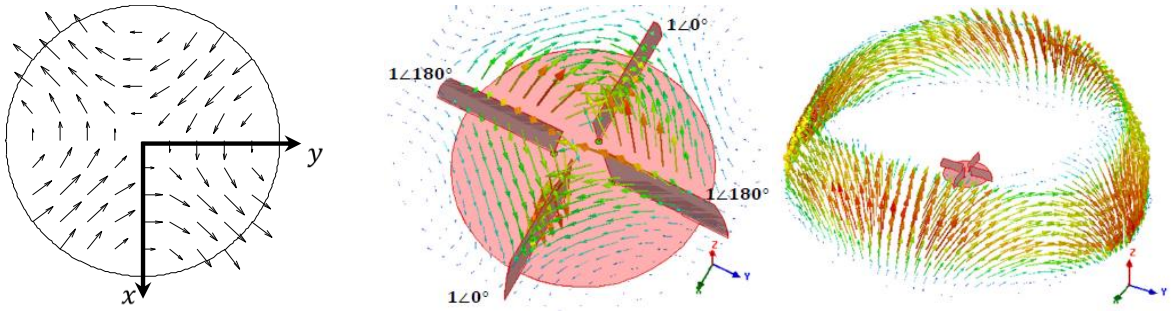
(a) (b) (c) (d)

Fig. 9. (a) Simulated S_{11} and maximum axial ratio in horizontal plane, (b) The average RHCP and LHCP gains in horizontal plane, (c) radiation efficiency, and (d) azimuth radiation pattern at 1.6 GHz.

bandwidth is from 1.11 GHz to 2.04 GHz (59.1%). The 3–dB bandwidth for the maximum value of axial ratio in all azimuth direction (horizontal plane) is from 1.21 GHz to 2.01 GHz (49.7%). Fig. 9(b) depicts the average RHCP and LHCP gain in all azimuth directions as a function of frequency. It can be seen that the RHCP gain is almost constant around 0 dBi for the entire band and the isolation between RHCP and LHCP gain is more than 20 dB on average for most frequencies and more than 15 dB for the entire frequency band. The radiation efficiency, which is shown in Fig. 9(c) is more than 80% for the entire band. The radiation pattern at azimuth plane is shown in Fig. 9(d) at 1.6 GHz. It shows good omnidirectional property in azimuth direction.

B. Low Profile Wideband Circularly Polarized Omnidirectional Antenna Based on Excitation of Two Orthogonal TE_{21} Modes

There are few designs for wideband CP omnidirectional antennas with more than 40% impedance and AR bandwidth. In this section, a novel low profile wideband CP omnidirectional antenna is introduced. The antenna operates based on excitation of a field distribution similar to circular TE_{21} mode in free space and in the near field of the antenna. Fig. 10(a) shows the electric field distribution of circular TE_{21} mode. Fig. 10(b) shows the structure for excitation a field distribution similar to Fig. 10(a) in free space. The structure is an array of four monopole sectorial loop antennas (SLAs)



(a) (b) (c)

Fig. 10. (a) Electric field distribution of circular TE_{21} mode. (b) Array of four monopole SLAs, excitation phases and amplitudes, and the resulting near field electric field. (c) Farfield electric field generated by the field distribution in (b).

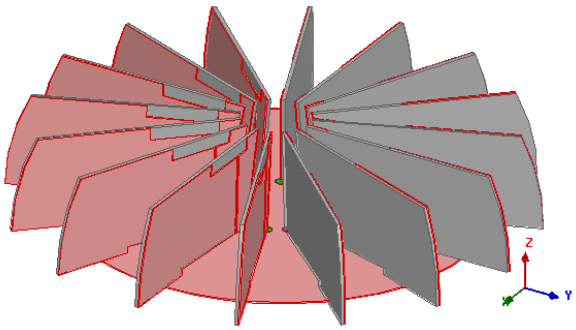
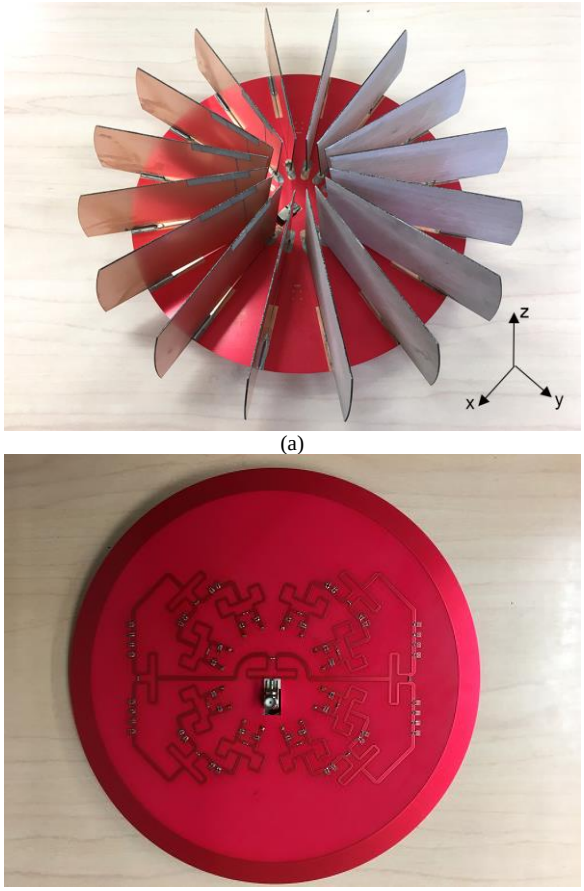


Fig. 11. Structure of the low profile wideband circularly polarized antenna.

mounted on a ground plane. SLA has been previously used for design of low profile wideband linearly polarized antennas. It can be shown that this field distribution generates a farfield linearly polarized electric field which its polarization tilt angle varies in azimuth direction. Fig. 10(c) shows the resulting farfield electric field generated by the field distribution in Fig. 10(b). Due to the symmetry in the field distribution of the mode, the direction of the linearly polarized electric field varies by 90° if azimuth angle increases by 45°, i.e. if in the azimuth angle φ the electric field is horizontally polarized, at $\varphi \pm 45°$, it is vertically polarized. This suggests that if two field distributions similar to circular TE_{21} mode are excited in free space such that one mode is spatially rotated by 45° with respect to the other mode, then in farfield and at each azimuth angle, two orthogonal linearly polarized electric fields are generated. In this case, omnidirectional circularly polarized radiation can be obtained by exciting the two modes by 90° phase difference. To have circular polarization operation, the minimum number of eight SLAs is required. However, simulation results show that with eight SLAs, the axial ratio bandwidth is 19%. To increase the bandwidth, number of SLAs is increased to 16 elements. Fig. 11 shows the antenna structure. To improve the radiation efficiency of the antenna at high frequencies, a gap is introduced in the SLA structure and also the height of the element is reduced in the locations close to the feed. The antenna can be fabricated using PCB substrates. Fig 12(a) shows the fabricated antenna. The diameter and total height of the antenna (excluding connector) are 250 mm and 47 mm, respectively. To excite the SLA elements by the required amplitudes and phases, a feed network consists of Wilkinson power dividers and blocks of low-pass/high-pass filters (providing constant 45° phase difference over the band 1.2–2



(a)

(b)

Fig. 12. (a) Fabricated CP antenna. (b) Feed network on the backside of the ground plane (for RHCP operation).

GHz) is designed. The feed is shown in Fig. 12(b) and fabricated on the backside of the ground plane. The feed in Fig. 12(b) is for RHCP operation.

Fig. 13(a) shows the measured reflection coefficient. It can be observed that the reflection coefficient is less than -10 dB over the band 1−2.1 GHz.

Fig. 13(b) shows the average of the measured axial ratio in horizontal plane. The antenna has the axial ratio less than 3-dB within the range 1.11–2.17 GHz (64% 3-dB axial ratio bandwidth).

Fig. 13(c) shows the average of the measured RHCP and LHCP gains and polarization isolation over the horizontal plane. It can be observed that within the frequency range 1.2–2 GHz, the RHCP gain varies in the range -2.5–0.5 dBi. The

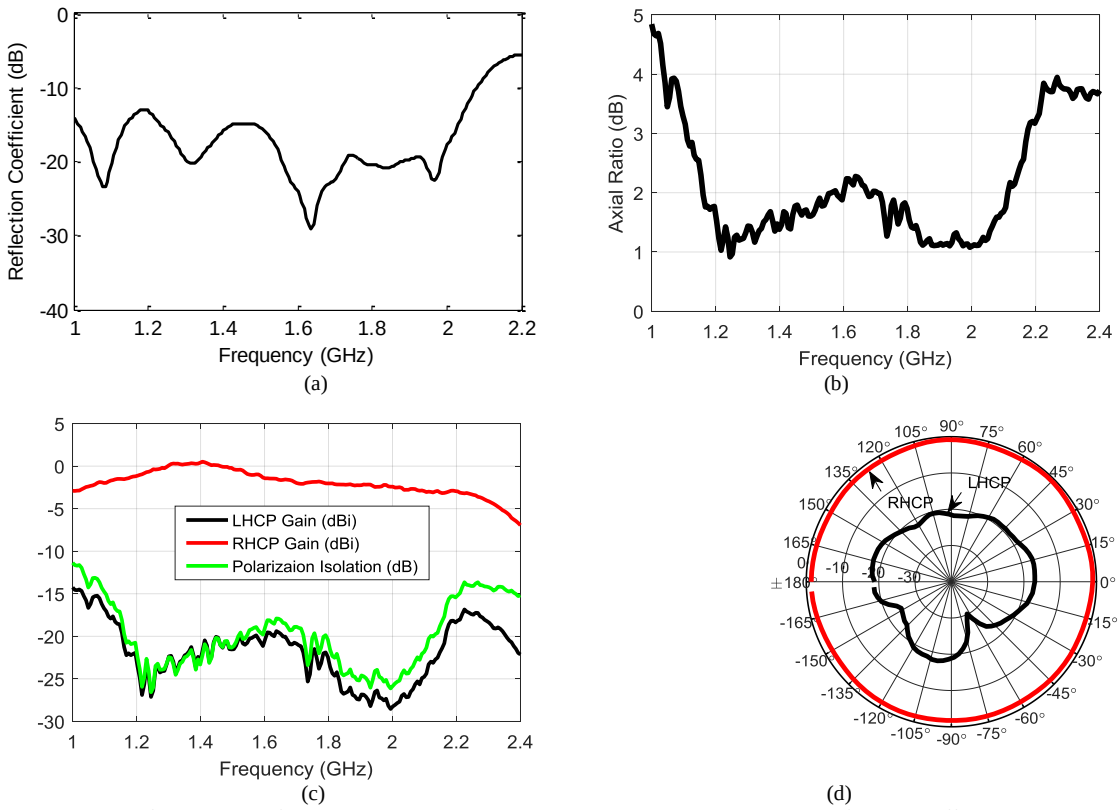


Fig. 13. Measurement results for the low profile wideband circularly polarized omnidirectional antenna. (a) Reflection coefficient. (b) Average of the axial ratio over the horizontal plane. (c) Average of CP gains and polarization isolation over the horizontal plane. (d) Gain pattern (dBi) in horizontal plane at 1.6 GHz.

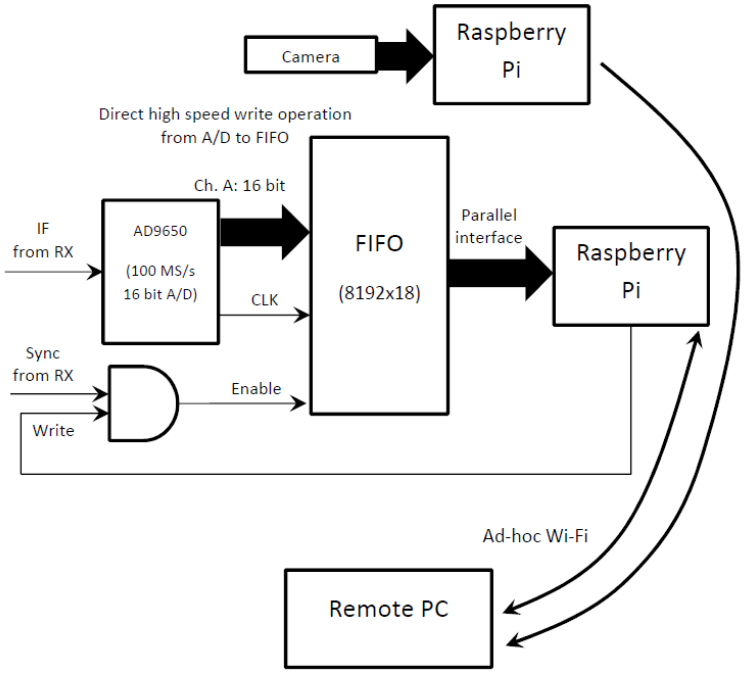


Fig. 14. Block diagram of the data acquisition module.

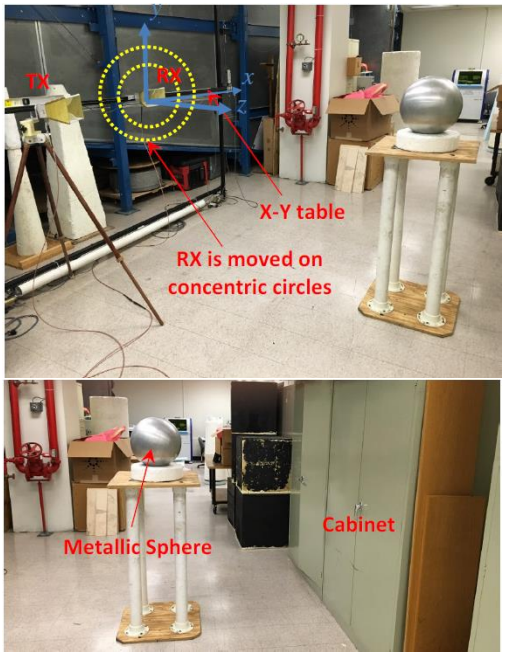
minimum measured polarization isolation within the specified frequency range is -18 dB.

Fig. 13(d) shows the measured gain pattern in horizontal plane at 1.6 GHz which indicates omnidirectional radiation pattern for the antenna.

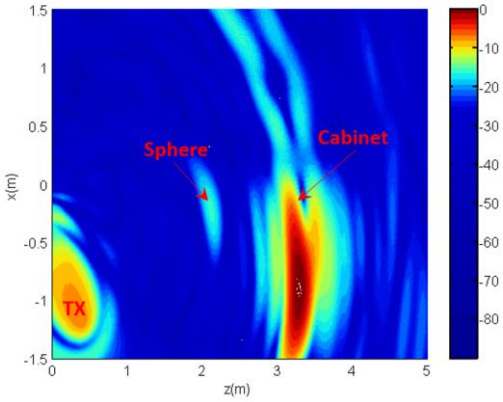
The minimum measured efficiency of the antenna is 68%.

V. A/D and Data Acquisition Module

Fig. 14 shows the block diagram of the data acquisition module for capturing the samples from high speed A/D and also the image from the camera. This system sends the acquired data to the main processing unit (a remote PC) via an ad-hoc Wi-Fi



(a)



(b)

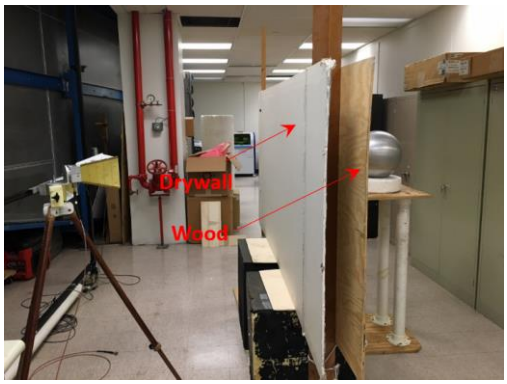
Fig. 15. (a) Measurement setup (no wall). (b) Imaging result for the configuration shown in (a).

network. This system allows direct access to the memory (FIFO) during the high speed (100 MSamples/s) write operation.

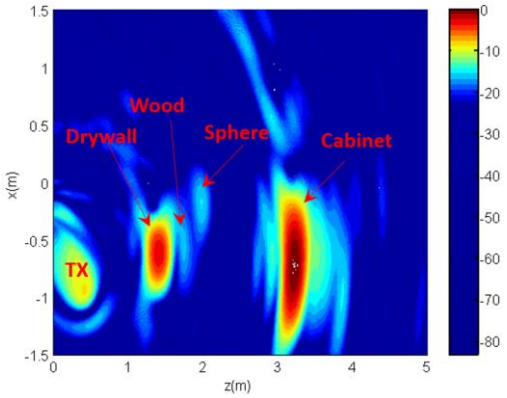
VI. Preliminary Measurement Results

To evaluate the performance of the proposed imaging method and also the fabricated transmitter and receiver module, a laboratory test has been performed using the fabricated radar. Fig. 15(a) shows the measurement setup. Transmitting and receiving antennas are two double ridged horn antennas. The position of the transmitting antenna is fixed whereas the receiving antenna is moved on concentric circles using an X-Y table as shown in Fig. 15(a). With this configuration for antennas, the measurement is bi-static. The object is a metallic sphere with diameter of 35 cm. In this measurement, a sample of the transmitted signal is mixed with the received signal (Rx and Tx are on the same board).

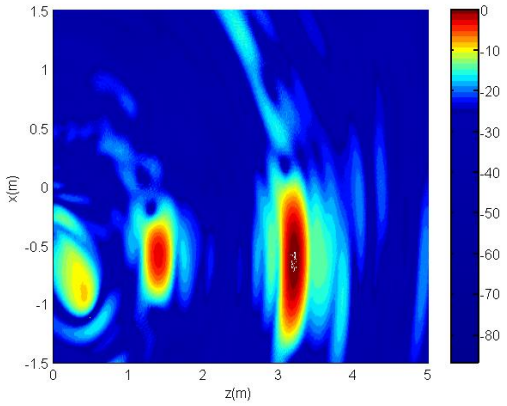
Fig. 15(b) shows the imaging result for the configuration shown in Fig. 15(a). The image is obtained by 51 samples of the



(a)



(b)



(c)

Fig. 16. (a) Measurement setup. (b) Imaging result for the configuration shown in (a). (c) Imaging result for the configuration shown in (a) while the sphere is removed.

scattered field. The sampling points are located on two concentric circles with the center located at the origin of the coordinate system and radii of 0.2 m and 0.4 m. The spacing between each two adjacent sampling points on a curve is $\lambda_{min}/2$. The measurement is performed over the frequency band 1.2−2 GHz. It can be observed in Fig. 15(b) that the sphere is clearly imaged. In addition to the sphere, the cabinet and the transmitter are also imaged. Image of the transmitter is due to the direct signal from the transmitter to the receiver.

Fig. 16(a) shows another measurement setup which includes a wall. The wall is formed by a layer of drywall and a layer of wood with separation of 20 cm. The metallic sphere is located behind the wall. The same configuration for transmitting and

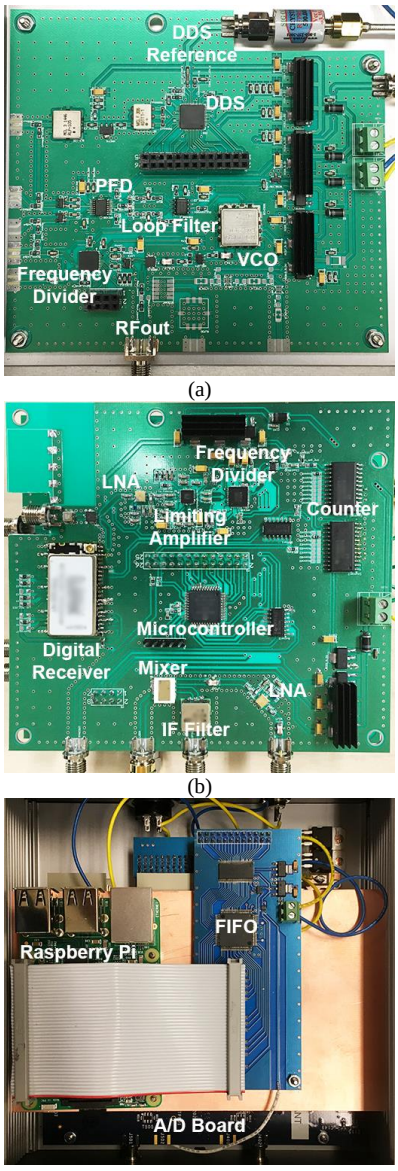


Fig. 17. (a) Transmitter/receiver circuit. (b) Synchronization and controlling circuit. (c) A/D and data acquisition module.

receiving antennas and sampling points as Fig. 15(a) is used here to form the image. Fig. 16(b) shows the imaging result. It can be observed that the sphere and also the wall boundaries are clearly imaged. Fig. 16(c) shows the imaging result while the sphere is removed. Comparing Fig. 16(c) and Fig. 16(b) indicates that the image of the sphere is removed in Fig. 16(c).

To test to the whole radar system, different parts of the radar were put in metallic boxes and power supply and overcurrent protection circuits were added to provide different voltage levels for different radar modules. Fig. 17(a) shows transmitter/receiver circuit inside the box. Fig. 17(b) shows the synchronization and controlling circuits inside the box. This circuit sits on the top of the transmitter/receiver circuit. Fig. 17(c) shows the A/D and data acquisition module.

Fig. 18(a) shows the configuration for the test of the whole radar system. In this configuration, there are three walls and one metallic sphere. Transmitter is located at the center and 143 sampling points are located on two circles with the centers

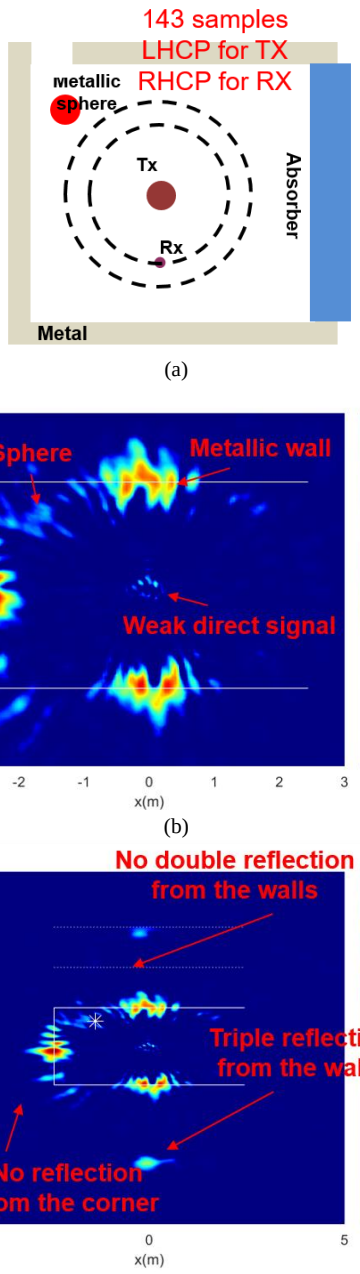


Fig. 18. (a) Imaging area and location of the Tx and sampling points. (b) Image. (c) Image in a larger domain, location of the double and triple reflection from the parallel walls on the top and bottom in (a) are indicated by dotted line.

located at the transmitter position and radii of 0.7 and 1 m. Fig. 18(b) shows the resulting image. In this image, the three metallic walls and the sphere are imaged. Also, it can be observed that the image of the transmitter at the center is weak. This is due to using orthogonal circular polarizations for transmitter and receivers (antenna introduced in part B of Section IV is used) which reduces the direct signal from the transmitter to the receiver. Fig. 18(c) shows the image formed in a larger domain. The location of the double and triple reflection from the parallel walls on the bottom and top are indicated by dotted lines. It can be observed that by using orthogonal circular polarization, the effect of double reflections for the parallel walls and corners are eliminated in the image.

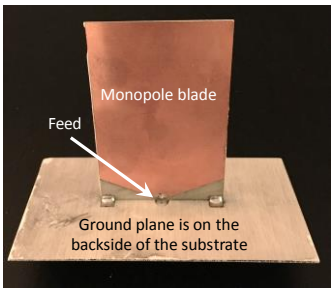


Fig. 19. Fabricated monopole blade antenna (wideband LP omnidirectional antenna).

The effect of the triple reflection from the walls in not removed since it has the same polarization as the signals which form the image of the objects and walls. Here, the image formed by triple reflection is significant since the walls are metallic and have large reflectivity. In practice, dielectric walls do not have large reflectivity and thus the effect of the triple reflections is negligible.

VII. Final Measurement Results

In this section, all-directions through-the-wall imaging method is tested using the fabricated radar system. To demonstrate imaging enhancement using CP antennas, the measurements have been performed using CP and LP antennas. The LP antenna used in measurements is a monopole blade antenna. The fabricated blade antenna is shown in Fig. 19. It is a wideband vertically polarized omnidirectional antenna with maximum gain of 2.1 dBi over the horizontal plane and in the range 1.2–2 GHz.

For the first measurement, a simple 3.6 m × 3.6 m × 1.2 m room was fabricated using two layers of drywalls and wooden posts. Fig. 20(a) shows the configuration and dimensions of the walls and the objects placed behind the walls. Objects are a rectangular corner reflector (with sides of 42.5 cm), two vertical metallic pipes (with diameters of 5 and 7.5 cm) connected from the sides, and one metallic missile launcher model placed horizontally. Fig. 20(b) shows the photograph of the imaging area. Transmitter is placed at the center of the room.

Fig. 20(c) shows the imaging result using LP antennas. According to Fig. 20(c), the walls, trihedral reflector, and pipes are imaged. The horizontally oriented launcher is not clearly imaged since it has low reflectivity for vertically polarized waves. The peak at the center of image is due to direct signal from transmitter to receiver.

To image the objects with large aspect ratio independent of their orientation, CP antennas can be used since they simultaneously transmit and receive vertically and horizontally polarized waves. In addition, using orthogonal circular polarizations for transmitting and receiving antennas, direct signal from transmitter to receiver is attenuated by polarization isolation of the CP antennas.

Fig. 20(d) shows the imaging result while left-handed CP antenna is used for transmitter and right-handed CP antenna is used for the receiver. It can be observed that the walls and all objects including the launcher is clearly imaged. In addition, no peak appeared at the center of image indicating that the effect of direct signal is reduced.

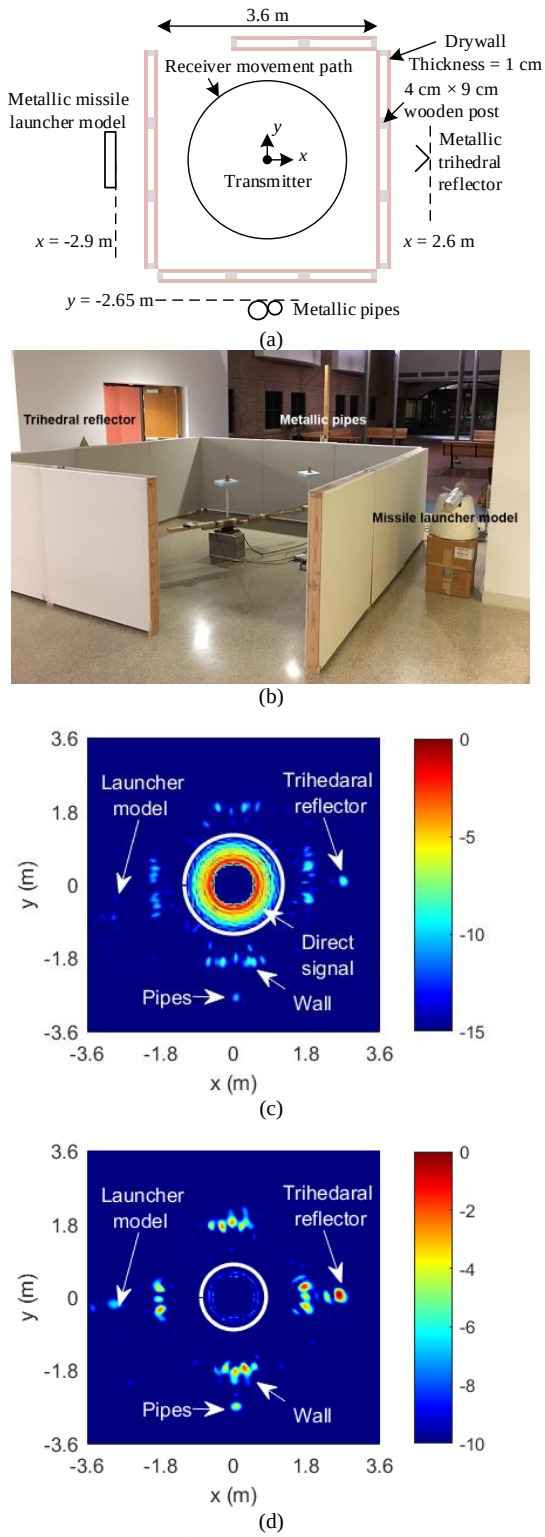
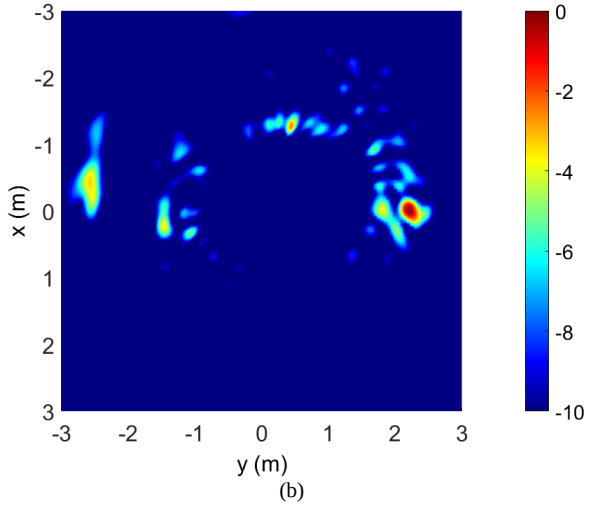


Fig. 20. (a) Geometry of the fabricated room and objects. (b) Photo of the room and objects. Normalized image (in dB) using LP antennas (c) and CP antennas (d). White circle shows the movement path of the receiver (181 sampling points). Transmitter is at (x, y) = (0,0).

It is noted that for image formation the effect of walls is not compensated and thus a small shift is observed in the position of the image of each object.



(a)



(b)

Fig. 21. (a) Measurement in a real room. (b) Normalized image (in dB) using CP antennas [181 sampling points, transmitter is at (x, y) = (0,0)].

Fig. 21(a) shows the second scenario which is a real room. With the same configuration for transmitter and sampling points as the first measurement, Fig. 21(b) shows the measurement results using CP antennas. It can be observed that the three walls shown in Fig. 21(a) are imaged. In addition, a wall behind the left wall in Fig. 21(a) is also imaged.

VIII. Conclusion

A brief review of the tasks completed before December 2015 has been presented. The tasks completed during December 2015 to December 2016 have been described in detail. The tasks are design, fabrication, and test of wireless synchronization circuits, wideband circularly polarized omnidirectional antennas, and A/D and data acquisition module. Samples of imaging using the fabricated radar have been also presented. The results show that the system can accurately image the features of the imaging area. Using orthogonal circular polarizations can significantly reduce the effect of the multiple reflections and direct signal from transmitter to receivers.

Publications

[1] B. Yektakhah and K. Sarabandi, "All-directions through-the-wall radar imaging using a small number of moving transceivers," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 11, pp. 6415-6428, Nov. 2016.

[2] X. Cai and K. Sarabandi, "A compact broadband horizontally polarized omnidirectional antenna using planar folded dipole elements," *IEEE Transactions on Antennas and Propagation*, vol .64, no. 2, pp. 414-422, Feb. 2016.

[3] B. Yektakhah and K. Sarabandi, "A wideband circularly polarized omnidirectional antenna based on excitation of two orthogonal circular TE_{21} modes," *IEEE Trans. Antennas and Propag.*, vol. 65, no. 8, pp. 3877–3888, Aug. 2017.

[4] Behzad Yektakhah and Kamal Sarabandi, "All-directions through the wall radar imaging using a synthetic array formed by a finite number of moving robots," *Antennas and Propagation Society International Symposium (APS/URSI)*, July 2014.

[5] Behzad Yektakhah and Kamal Sarabandi, "All-directions through the wall radar imaging enhancement using orthogonal polarizations and generalized pencil of function method," *Antennas and Propagation Society International Symposium (APS/URSI)*, July 2015.

[6] X. Cai and K. Sarabandi, "A compact broadband horizontally polarized omnidirectional antenna using folded dipoles," *Antennas and Propagation Society International Symposium (APS/URSI)*, July 2015.

[7] B. Yektakhah and K. Sarabandi, "A low profile wideband circularly polarized omnidirectional antenna using an array of sectorial loop antennas," *IEEE Antennas and Propagation Society International Symposium (APS/URSI)*, 2016.

[8] X. Cai and K. Sarabandi, "Broadband omni-directional circularly polarized antenna based on vertically and horizontally polarized elements," *IEEE Antennas and Propagation Society International Symposium (APS/URSI)*, 2016.

[9] B. Yektakhah and K. Sarabandi, "All directions through the wall imaging using omnidirectional bi-static FMCW transceivers," *IEEE International Symposium on Antennas and Propagation & USNC/URSI National Radio Science Meeting*, San Diego, CA, USA, 2017, pp. 233-234.

[10] Behzad Yektakhah and Kamal Sarabandi, "All-directions through the wall imaging using a small number of moving omnidirectional transceivers," *IEEE International Geoscience and Remote Sensing Symposium*, July 2017.

2017 Final Report Metrics

Task Number: SSP-17-2.3

Metrics

Notes: *Please list metrics for the reporting quarter only and only Publications that are accepted and searchable on web.*

Awards and/or Recognitions this Quarter

Honorable Mention: B. Yektakhah and K. Sarabandi, "All directions through the wall imaging using omnidirectional bi-static FMCW transceivers," *IEEE International Symposium on Antennas and Propagation & USNC/URSI National Radio Science Meeting*, San Diego, CA, USA, 2017.
http://2017apsursi.org/StudentPaperCompetition.asp

Refereed/Journal Papers accepted this quarter

[1] B. Yektakhah and K. Sarabandi, "A wideband circularly polarized omnidirectional antenna based on excitation of two orthogonal circular TE_{21} modes," *IEEE Transactions on Antennas and Propagation*, vol. 65, no. 8, pp. 3877–3888, Aug. 2017.

Conference Presentations/proceedings presented this quarter

[1] B. Yektakhah and K. Sarabandi, "All directions through the wall imaging using omnidirectional bi-static FMCW transceivers," *IEEE International Symposium on Antennas and Propagation & USNC/URSI National Radio Science Meeting*, San Diego, CA, USA, 2017, pp. 233-234.
[2] Behzad Yektakhah and Kamal Sarabandi, "All-directions through the wall imaging using a small number of moving omnidirectional transceivers," *IEEE International Geoscience and Remote Sensing Symposium*, July 2017.

Joint Papers with ARL presented or accepted this quarter – duplication with journals/ conferences expected. – n/a

Invention Disclosures for this quarter – n/a

Patent Applications Filed this quarter – n/a

Patents Awarded this quarter – n/a

PhDs Award this quarter: - n/a

Masters Awarded this quarter – n/a

Students supported fully or partially (by name)

MS: *Name(s)*
PhD: *Behzad Yektakhah*
Post Docs: *Name(s*
Master Degree: *Name(s*

PhDs Awarded: *Name(s*

MAST CTA Member Rotations into ARL

Name(s

ARL Researchers Working at Member Sites

Name

Workshops and Seminars conducted:

Title

MAST Quarterly Metrics Summary Table (Numbers in this table should correspond to the list above):

MAST Final REPORT METRICS TABLE (The data entered into this table should be for the reporting quarter only)	
Awards (Miscellaneous Recognitions and Awards)	1
Refereed/Journal papers	1
Conference presentations/proceedings	2
Joint Papers with ARL	0
Invention disclosures	0
Patent applications filed	0
Patents awarded	0
PhDs Award	0
Masters Awarded	0
Post Docs Supported	0
Master Degree Students Supported	0
Undergraduate Students Supported	0
CTA Member rotations into ARL	0
Visiting Researchers from ARL to the Member Sites	0
Employees hired by ARL from CTA students supported	0

Air Ground Mapping of Feature Rich 3-D Indoor Environments with RGBD Sensors

SPP 16-17.2.5

Camillo J. Taylor, Giuseppe Loianno, Armon Shariati, and Anthony Cowley
University of Pennsylvania
GRASP Lab, Levine Hall Rm 474
3330 Walnut St, Philadelphia, PA 19104
Phone: 215-898-0376
{cjtaylor,loiannog,armon,acowley}@seas.upenn.edu.com

M. Ani Hsieh, Daniel C. Mox, and Hadi Hajieghrary
Drexel University
Scalable Autonomous Systems Lab
Mechanical Engineering & Mechanics, Randell 115
3141 Chestnut St, Philadelphia, PA 19104
Phone: 215-895-5870
{mhsieh1, dcm64,hh449}@drexel.edu

***Abstract* — We describe the development of MAST-scale platforms and software leveraging recent attempts to commercialize color/depth cameras, and develop autonomous platforms that are capable of three-dimensional situational awareness. Specifically we report on our developments of micro unmanned aerial vehicles equipped with RGB-D sensors and camera phones capable of on-board localization, mapping, and navigation; construction of high fidelity 3D maps using depth panoramas; a multi-robot information based exploration strategy that guides the robots to best locations to obtain these depth panoramas using a novel angle occupancy grid; and the ability to recover the relative position and orientation of a pair of moving platforms by fusing estimates provided by visual odometry and distance measurements provided by ultrawide band ranging radios. Extensive experimental and simulation studies are presented to highlight these accomplishments**

***Index Terms*—Component, formatting, style, styling, insert.** *(key words)*

I. Introduction

Current operational scenarios envision a combination of man-throwable robots and man-portable systems being used to create an accurate map of the environment for search and rescue operations. This project addresses the development of MAST-scale platforms and software leveraging recent attempts to commercialize color/depth cameras, and develop autonomous platforms that are capable of three-dimensional situational awareness. Current work in MAST is based on laser scanners and cameras. Laser scanners are not well suited to SWAP-constrained platforms. While cameras are compatible with MAST-scale platforms, their ability to generate three-dimensional models is somewhat constrained by illumination conditions and inability to directly provide scale information. RGB-D cameras have the potential to add new capabilities to fast moving quadrotors and ground robots, particularly in constrained, indoor environments.

Our main objective is to develop a networked team of aerial and ground robots that will operate as a cohesive group in an indoor environment with the ability to autonomously navigate and create three-dimensional maps. Our aerial platform is a quadrotor (AscTec Hummingbird scale hex rotor aircraft) equipped with RGB-D sensor (Primesense or Google Tango prototype) and an IMU. Our mobile ground robots are small mobile differential drive platforms equipped with similar sensors. An aerial robot will be used to localize the ground robot team and to map the environment. Ground robots will be used to refine areas on the map that are unreachable by the aerial robot. Moreover, due to their small size, they will be equipped with specialized sensors that enable them to enter narrow confined areas and obtain information from unique vantage points. A secondary objective is to develop mutual information based active control and exploration schemes that determine trajectories that maximize the mutual information of the map with the sensory information gleaned from traversing the trajectory. Our framework will allow for heterogeneous sensors and heterogeneous platforms in three-dimensional settings.

This project will enable the deployment of aerial and ground robots in an indoor environment providing soldiers with the information about the layout of the environment. The map can be employed to identify humans, booby-traps elements, as well as objects, which can represent potential threats. Furthermore, as ground and aerial vehicles traverse the environment, they should have the ability to discern potential threats using their onboard sensors when rounding a corner or entering a room. The role of both robot types is complementary. While the aerial platform navigates through the space, the ground robots would refine the map focusing on regions inaccessible to the aerial platform. The final objective is to enable localization of one or more targets located inside buildings that may not be immediately visible to a single vehicle or soldier (e.g., a backpack under a table) but is visible to a heterogeneous team of robots. Operations, such as landing, grasping and perching, would be beneficial to interact with localized targets.

There are a number of Simultaneous Localization and Mapping (SLAM) approaches for Micro Aerial Vehicles

(MAVs). Good results have been obtained using monocular cameras and IMUs (Inertial Measurements Units) [1], stereo camera configurations [2] and RGB-D sensor systems [3], [4], [5], [6], [7], [8]. In [3], [4], the Kinect depth data and on-board vehicle sensors are used to perform state estimation through Kalman filtering, while in [1], the same filter is used to combine monocular visual information with inertial sensor data solving the 3-D SLAM problem on a computationally-constrained aerial platform. In particular, in [5] an RGB-D 3 D mapping system utilizes a novel joint optimization algorithm combining visual features and shape-based alignment. In [7] a direct 3-D tracking approach is proposed such that an error is based directly on the intensity of pixels. Other algorithms are based on fusing depth maps to a coherent 3-D model [8].

Most previous approaches are designed for use on a single platform focusing only on mapping, rather than localization, and are too slow to be used onboard MAVs. Moreover, in some cases, to map large environments it may be necessary to deploy additional vehicles to cooperatively accomplish the mapping task. Through collaboration, multiple MAVs can more quickly and reliably cover the same environment as compared to a single vehicle. Moreover, every vehicle can exploit the fused map information when making decisions on next steps and task allocations.

Since different vehicles can be launched from different initial positions and orientations in a collaborative mapping task, we cannot assume any prior knowledge on the relative pose of the vehicles. Thus, the relative pose information should be inferred from different measurements (that may or may not be independent) of the external scene by different vehicles. This problem is addressed in [9] where two MAVs map the same environment. Loop closure detection between different cameras, map merging, and concurrent map access are performed on a ground station where the maps are replicated. The advantage of this approach is that each vehicle starts its own map. However, the ground station is responsible for map merging and must manage redundant information like the maps built by the individual vehicles, including any overlap in the mapping data.

In general, mapping strategies that rely solely on monocular vision are often not robust to sudden changes in depth. While there are mapping strategies developed for ground robots equipped with range sensors in 2-D environments [10], [11], [12], [13], lasers are generally ill- suited for SWAP-constrained platforms and cannot provide complete information of the 3-D environment. The use of bearing measurements from monocular cameras coupled with IMU has been exploited in [10]. An interesting approach is presented in [15] where 3 synchronous cameras are able to perform localization even in dynamic environments. Since the images must be synchronized in this approach, the strategy is hard to extend to MAVs.

In [16], a stochastic approach is presented for cooperative mapping with a Kalman Filter fusing different camera poses and observed landmarks. The combination of synchronized data between the RGB and the depth sensors for multiple robots can be exploited to create a complete framework for localization and mapping. The central idea is to decompose the problem into (a) a centralized monocular SLAM problem with sparse representation involving features that are tracked by individual vehicles avoiding single cameras scaling problems and map merging scaling issues like in [9]; and (b) the problem of associating robot poses and depths with features to creates a dense 3-D map, a problem that can be solved in a distributed way. By decomposing the problem in this way, we avoid the computational bottleneck of 3-D RGB-D cooperative SLAM and achieve increased robustness to noise in depth, which is typical in outdoor or brightly lit environments.

The research discussed in this task addresses questions related to cooperative localization and mapping for autonomous operation in complex and challenging environments. In particular, we focus on addressing the key limitations of methods in the state of the art, by using fast visual algorithms for localization, high-resolution dense reconstruction and by employing cooperative strategies with multiple vehicles. Specifically, our contributions include:

- The development of autonomous micro aerial vehicles equipped with RGB-D sensors and camera phones capable of autonomous, on-board localization, and providing a sparse and dense map of the environment. The vehicles use their onboard IMU and RGB-D sensor to navigate autonomously.
- The use of depth panoramas in the construction of detailed 3D models of extended environments. We show how a properly equipped robot can acquire these panoramic images by collecting sequences of depth images with a commodity depth sensor. These sequences are stitched into RGBD panoramas that efficiently capture scene information while reducing noise in the captured imagery. Scene structure is then extracted from the RGBD panoramas and used to register a collection of panoramas to obtain coverage of extended areas. The strategy maintains fine geometry detail in 3D reconstructions, uses a modest amount of memory, and enjoys excellent scaling characteristics across large environments.
- A multi-robot information based exploration strategy that maximizes the Cauchy-Schwarz Quadratic Mutual Information (CSQMI) on a novel angle enhanced occupancy grid to guide robots in the collection of RGBD panoramas. The angle enhanced occupancy grid emphasizes perspective in addition to coverage which yields higher quality maps when compared to traditional occupancy grid methods. Compared to frontier exploration, the angle enhanced information maximizing strategy yields better coverage of the space from various view angles.
- An approach to recover the relative position and orientation of a pair of moving platforms by fusing the kinds of estimates provided by visual odometry systems with distance measurements obtained from ranging radio systems. Combining these two complementary sources of information provide estimates that could not be obtained using either source separately. In addition, we show that the

localization problem can be posed as a convex optimization problem where the structure of the problem is leveraged to provide an efficient and effective solution. Experimental results demonstrate that the proposed method performs creditably on actual robotic platforms.

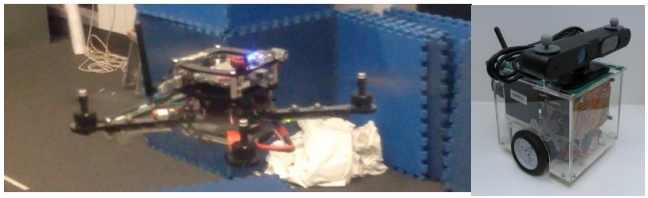


Fig. 1. A Hummingbird quadrotor (left) and micro UGV platform (right) equipped with RGB-D sensors.

II. Related Work

- Synergy with Existing Works and External Collaborations
- Technology Transitions summary

This work builds upon Penn's preliminary work on developing an autonomous navigation system for a micro UAV for operation in indoor environments that relies on RGB-D sensor localization. This system consists of an Asctec Hummingbird quadrotor equipped with a small computational platform (Fig. 1). It fuses information from an onboard RGB-D sensor and IMU to perform localization, plan and control without any operator intervention or any external localization system (like GPS or Vicon). A complementary approach, based on GPU, useful to obtain a dense reconstruction with modest memory occupancy requirements, has also been developed and tested on a consumer laptop. Drexel has developed a fleet of RGB-D equipped small-scale ground robots capable of exploring confined environments, providing vehicle odometry and a map of the scanned environment.

Penn has also developed an algorithm that allows Primesense RGB-D cameras on quadrotors to generate 3-D maps, supporting localization and mapping at 30 Hz on a limited computational platform. This can, in principle, also be used for ground robots and for air ground coordination. The algorithm is able to generate a sparse map running in parallel with respect to a dense one obtained from the IR projector. The user has the possibility to choose the type of environment representation according to the employed platform and to his needs.

A complementary approach to the strategy mentioned previously, , with respect to the one mentioned and GPU based, for high-resolution dense environment reconstruction is presented in the next section. The use of this sensor is motivated by its low weight and by the possibility to use the depth data to solve the scaling problem affecting monocular SLAM, used in this context by the RGB camera to solve the localization task. Finally, preliminary experiments on cooperative mapping between two aerial platforms, based on collected data, running on a ground platform, have been recently performed (Fig. 2).

The research topics in this project range from cooperative localization and mapping for autonomous operation in complex and challenging environments. In particular, the work focuses on addressing the key limitations of methods in the state of the art, by using fast visual algorithms for localization, high-resolution dense reconstruction, and by employing cooperative strategies with multiple vehicles. The techniques, developed through this thrust, are highly relevant as tools in an autonomous aerial vehicle's repertoire and will be transitioned to the ARL aerial vehicles

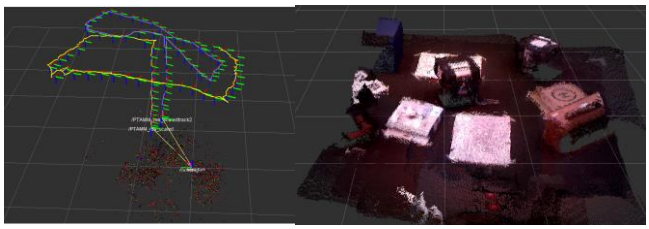


Fig. 2. Cooperative mapping with two robots, each equipped with a RGB-D camera (left), localizing relative to each other creating a sparse map (left), while creating a dense map (right).

III. Approach

A. Real Time Dense RGB-D Mapping and Alignment

In another thrust, we are developing a high-performance system around a novel sparse 3D geometry representation. The aim is to enable high-frame rate capture of 3D data over very large physical spaces with modest hardware requirements. The system uses a loosely coupled, modular world representation that can be easily replaced to compare different methods.

The software takes advantage of multiple CPU cores and GPU processing by building upon composable streaming primitives that offer high throughput, low latency, and low memory cost. The streaming network is also amenable to rich output monitoring. Not only are camera pose estimates and the world representation continuously updated, telemetry and visualization output may be attached to various points in the system to provide a detailed view of the internal working of the system during development and tuning.

The current version of the software features traditional ICP pose registration techniques and dense geometry representations. The attached images show the live output of the system reconstructing a 32 cubic meter environment at 25 Hz using the integrated GPU on an ultra-portable laptop. The data set shown here is from the paper, "Dense Scene Reconstruction with Points of Interest" by Qian-Yi Zhou and Vladlen Koltun (SIGGRAPH 2013). The images (see Fig. 3) show the live 3D reconstruction in the lower-left, and the raw range data played back from a log file in the upper-right. The attached chart (see Fig. 4) shows performance telemetry that breaks down computation time between the major cost centers of the system for each frame, in the upper left, and averaged across the entire data set, in the lower right.

Our work in this area is motivated by the observation that 2D depth maps can provide compact, high-resolution representations of 3D environments. We plan to explore

methods where extended areas are represented by a collection of depth frames linked by relative pose estimates. Such a representation would be well suited for the envisioned scenarios where multiple platforms would contribute views or frames to a global representation. We further intend to develop algorithms that would allow the UAVs and UGVs to both plan and execute collision free paths in highly cluttered environments using these topometric representations.

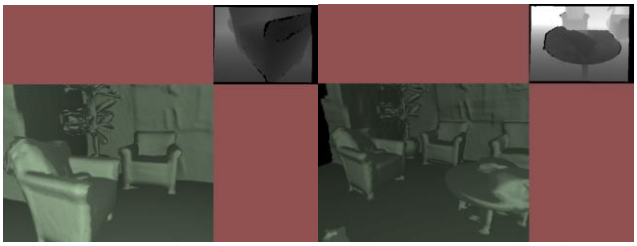


Fig. 3. Examples of models produced by the current version of our real-time RGB-D reconstruction algorithm from the RGB-D data set collected by Koltun et al.

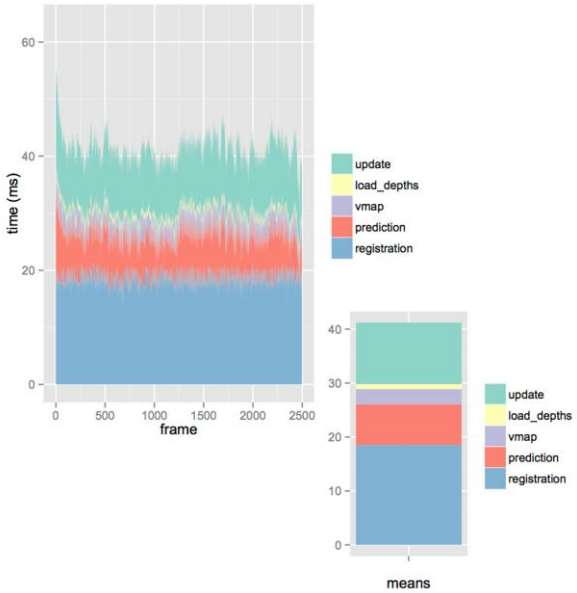


Fig. 4. Timing breakdown of the stages in our real time RGB-D analysis system on an ultraportable laptop.

B. State Estimation and Control

We plan to use the RGB-D sensory information along with IMU information with a UKF to generate pose estimation at 100 Hz to perform autonomous onboard control and navigation. The actual configuration will run on an Odroid embedded platform mentioned in Table I.

Let us consider a set of 3D points in the IR camera frame

$$\mathbf{P}^{IR} = \left[\ \mathbf{p}_1^{IR}, \cdots, \mathbf{p}_n^{IR} \ \right]^{\mathrm{T}}, \tag{1}$$

where is a 3D point expressed in the IR frame. The 3D points can be expressed in the current RGB frame using the extrinsic camera calibration parameters

$$\mathbf{p}_i^{RGB} = \mathbf{t}_{IR}^{RGB} + \mathbf{R}_{IR}^{RGB} \mathbf{p}_i^{IR}, \tag{2}$$

A set of scale factor can be computed in the current frame comparing the Euclidean norm of each monocular SLAM 3-D point with the corresponding depth data

$$\Delta p_i^{RGB} = \|\mathbf{p}_i^{RGB}\| \tag{3}$$

$$\Delta p_{v,i} = \|\mathbf{p}_{v,i}\| \tag{4}$$

$$\mathbf{s} = \left[\ s_1, \cdots, s_n \ \right]^{\mathrm{T}}, \ s_i = \frac{\Delta p_i^{RGB}}{\Delta p_{v,i}} \tag{5}$$

A statistical procedure is then considered to delete the scale values, which are outside twice the standard deviation range. The scale factor is computed when the bundle adjustment is converged since in this way, the scale drift can be avoided. This approach avoids an extra constraint in the bundle adjustment procedure, which will obviously affect the computational performance. However, under investigation is the possibility to employ depth filters [19] to refine the map and employ a faster monocular algorithm [20] instead of the one used in this context. Due to the algorithm speed then, it will be probably possible to insert as an explicit constraint the direct depth measurement in the bundle adjustment. Once the scale factor problem is solved, the monocular SLAM pose estimation can be corrected and the full vehicle's pose is estimated. The visual feedback is provided at 30 Hz.

To enable onboard control, a UKF estimator is used to fuse vision and inertial data. This process enables us to estimate the full pose at 100 Hz. The state is represented by

$$\mathbf{x} = \ [x \ , \ v \ , \ \Phi \ , \ \mathbf{a}_b]^T \ , \tag{6}$$

where the quaternion is represented by the vector

$$\Phi = \ [q_0 \ , \ q_1 \ , \ q_2, \ q_3] \ , \tag{7}$$

and the accelerometer biases by

$$\mathbf{a}_b = \ \left[a_{b_x} \ , \ a_{b_y} \ , \ a_{b_z}\right] \ . \tag{8}$$

In this case, vehicle's orientation is represented by quaternions.

The prediction step consider linear acceleration and angular velocities measurements given by the IMU

$$\begin{aligned} \mathbf{u} = \left[\ \omega, \mathbf{a} \ \right]^{\mathrm{T}} = \left[\ \omega_x, \omega_y, \omega_z, a_x, a_y, a_z, \ \right]^{\mathrm{T}}, \\ \mathbf{v} = \left[\ \mathbf{v}_\omega, \mathbf{v}_\mathbf{a}, \mathbf{v}_{\mathbf{a_b}} \ \right]^{\mathrm{T}}, \\ \mathbf{x}_{k+1} = f\left(\mathbf{x}_k, \mathbf{u}_k, \mathbf{v}_k\right). \end{aligned} \tag{9}$$

where $\mathbf{v}_k$ represents Gaussian noise. Finally, the vision data is used update the output linear measurement model

$$\mathbf{z}_k = H\mathbf{x}_k + \eta_k, \tag{10}$$

where η_k is a Gaussian noise. The pose information is then used by a SE(3) controller [17], [18] to control and plan the motion of the vehicle in the unknown unstructured environment.

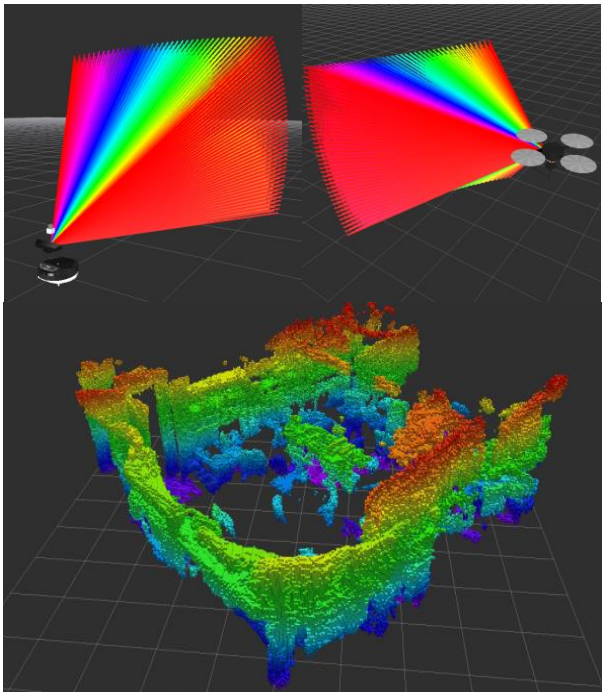


Fig. 5. Sensor model for a ground robot (top left), aerial robot (top right) and a dense depth map of an indoor environment (bottom).

C. Mapping and Localization

Our initial approach will merge the data obtained from one aerial vehicle and multiple ground robots. The aerial robot can localize ground robots using vision-based localization algorithms. The final exploration task aims to provide information in the scanned environment of specific areas and to identify potential dangerous objects such as backpacks. Once the overall system has performed the identification, the aerial or a ground vehicle should, eventually, perform operations like perching, grasping, pushing and interacting with potential dangerous objects. The ground station is just used to merge the information provided by the different platforms and optimize the alignment information to obtain a coherent 3D map, which can be used to further refine the object position as well as landing and grasping areas. The main approach, to solve the cooperative localization task, consists of one aerial vehicle starting the map and any another one relocalizing on it, through shared visual data information. Once relocalization is completed, the new vehicle can contribute to the map' s expansion. This approach can be employed to localize different ground vehicles. Moreover, it can be employed to solve the aerial ground relative localization task. However, in this last case, to deal with the difficulty of identifying, from large cameras viewpoints changes overlapping regions, this problem can eventually be solved observing uniquely blinking LED patterns on each ground from the aerial perspective. Redundant identification information can be conveniently merged and optimized using a consensus and graph based algorithm on the ground station.

All robots send to a ground station the map information where it is processed and refined.

D. Strategy for Exploration

We plan to pursue a model predictive control algorithm where the cost functional is based on information. We know that Shannon and Renyi entropy offer us metrics of measuring uncertainty. However, both metrics focus on precision and quantifying the bits of uncertainty on information rather than explicitly encoding actionable information. In our approach, we will formulate a cost functional that will encode actionable information. We will define the mutual information between the environment map, m, and observations, z, that are likely to be made while executing a set of control actions, that will encode the mutual information. We then propose a cost functional that encodes the mutual information gained by executing a trajectory, and find the trajectory that maximizes this cost functional

$$\arg\min_{x\in\mathcal{X}} I[m, z_\tau(x_\tau)] \quad . \tag{11}$$

Note that the cost functional requires a double integration or summation over the possible environments and the possible measurements that can be made. As such this computation grows prohibitively with the increase in the environment complexity and the richness of the sensor measurements. To address this we will pursue analytical approximations to the expression to guarantee performance with sub optimality bounds while allowing real-time performance.

It should be noted that this approach naturally allows us to deal with different sensor models and different control actions, and lends itself to exploration with heterogeneous teams.

E. Communication

In multirobot estimation and control, the questions of who knows what and when, and who makes what decision and when are quite central. We will pursue efficient sparse, local representations of the environment enabling robots to exchange information in real time to enable distributed decision making while ensuring that all robots have the same sparse map. Instead, the dense reconstruction, which is done locally, will not be available across the robot team. However, the reconstructed dense map can be centrally aggregated.

F. Platform and sensors

Based on algorithms proposed, we want to use suitable hardware that is capable to fulfill the previous mentioned requirements. We propose a quadrotor and hexarotor platform packages. They are equipped with a RGB-D sensor and Google Tango device running onboard vision-based SLAM. In case it is needed, we will design a lightweight, 5 gram carbon fiber shell for both platforms, so that they are resilient to collisions to the wall without causing failure to the propeller blades. The processor unit will be an Odroid platform. See Table I for configuration details. Both platforms will be supported by an ARO DURIP grant, awarded to Dr. Kumar and developed in collaboration with Dr. Loianno.

TABLE I: CONFIGURATION TO BE EXPLORED FOR VISION BASED LOCALIZATION AND MAPPING

	Package
Processor unit	ODROID-U3, ODROID-XU and ODROID-XUE
Camera	RGB-D sensors and Google Tango
Platform	KMel NanoHex, Asctec Hummingbird and COTS vehicles

IV. RESULTS

A. State Estimation, Control, Mapping and Localization

In this work, we have developed autonomous micro aerial vehicles based on RGB-D sensors and camera phones. The vehicles equipped with RGB-D sensors are able to localize autonomously and providing a sparse and dense map of the environment. Moreover a ground station is responsible for map merging when the two vehicles observe a common area. The vehicles use their onboard IMU and RGB-D sensor to navigate autonomously. The control, localization and planning is running on a computationally limited platform such as the Odroid-XU3.

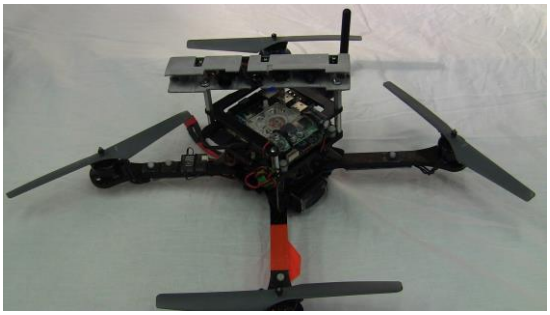


Fig. 6. The aerial vehicle based on RGB-D sensor

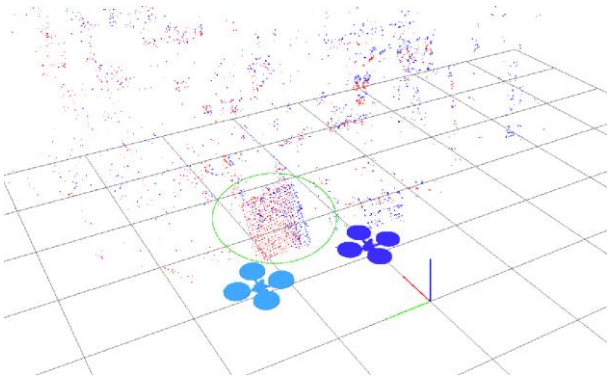


Fig. 7. Overlap detection and map merging.

Since the Odroid-XU3 has the same processor as an off-the shelf camera phone, we started to investigate the possibility to provide autonomy to these vehicles using a simple phone. Initial tests have successfully established the possibility to provide full autonomy to a Micro Aerial Vehicle using a specialized phone like the Google Tango and a small Odroid board. However, recently we have shown the possibility to use only an off-the shelf phone such as a Samsung Galaxy S5. The planning, control and estimation is running at 200 Hz and can be installed as an app potentially on any Android phone.



Fig. 8. The camera phone based aerial vehicle design.

Furthermore, we showed the feasibility of cooperative reconstruction, inter-robot collision avoidance and cooperative mapping using multiple vehicles carrying smartphone devices. All the control, planning and estimation tasks are solved on each robot in a single app using the camera and IMU in the phone. The architecture, which enables flight of multiple aerial platforms powered by smartphones, is mostly distributed. Each vehicle is able to perform state estimation, planning and control tasks. The only role of the ground station is to solve the goal assignment problem and to collect data from each platform such to build a cooperative map of the environment. The focus now is to bring this technology on smaller platforms. The vehicles are able to fly autonomously, avoid collisions between robots and reconstruct the environment in a cooperative manner. Our methodology present as the only drawback, the necessity to fly in a free obstacle space environment. We are currently investigating how to solve this problem and make the method more general.

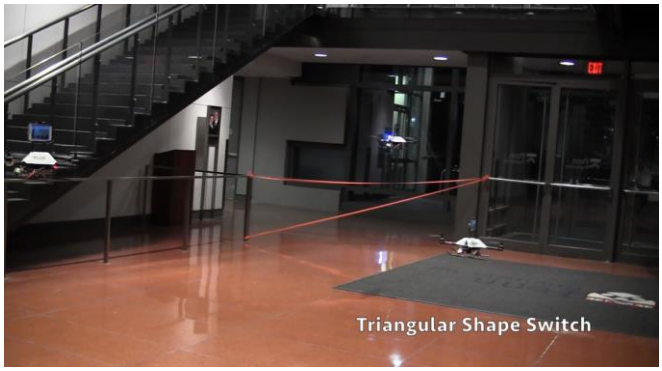


Fig. 9. Formation navigation with multiple vehicles.

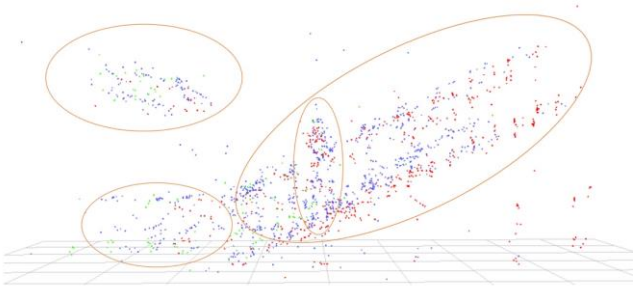


Fig. 10. Cooperative reconstruction from multiple vehicles.

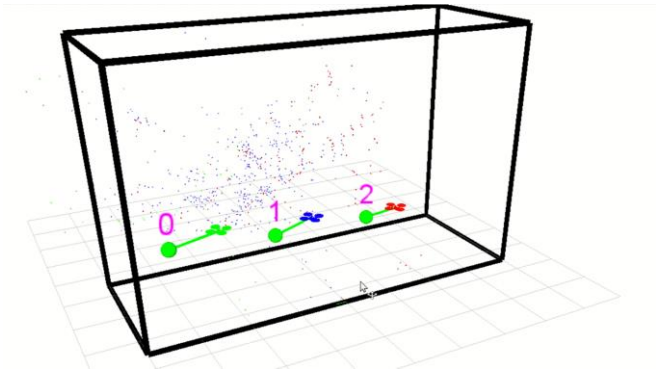


Fig. 11. User interface with goal selections and reconstructions.

B. Real Time Dense RGB-D Mapping and Alignment

Our work in this area is motivated by the observation that 2D depth maps can provide compact, high-resolution representations of 3D environments. We are interested in methods where extended areas are represented by a collection of depth frames linked by relative pose estimates. In our work, we demonstrated the use of depth panoramas in the construction of detailed 3D models of extended environments. We have developed an approach for the acquisition of such panoramic images using a robotic platform that collects sequences of depth images with a commodity depth sensor. These sequences are stitched into panoramic images that efficiently capture scene information while reducing noise in the captured imagery. Scene structure is extracted from the panoramas and used to register a collection of panoramas to obtain coverage of extended areas. The presented approach maintains fine geometry detail in 3D reconstructions, uses a modest amount of memory, and enjoys excellent scaling characteristics across large environments.

To acquire the requisite measurements for our depth panorama, we equipped a MAST proxy ground vehicle with a Kinect One RGB-D as shown in Fig. 12. Using this system, we systematically collected depth frames as the robot rotated and translated through the scene. We then correlated the RGB-D images with the odometry measurements provided by the vehicle.

Fig. 13 shows an example of an output produced by our novel RGB-D panorama algorithm. The algorithm achieves this by gauging the transformation between successive frames by leveraging correspondences. A global optimization scheme was employed to handle the loop closure and produce a consistent final panoramic image.

More recent work in this area has allowed us to extend this technique to incorporate RGB imagery. This involved a careful calibration of our RGB-D sensor to recover the relative position and orientation of the depth sensor and the color camera. Once this had been done we were able to transfer depth measurements from the depth sensor frame to the color frame. By incorporating the RGB imagery we were able to leverage feature correspondences found in the image to further improve and refine the positions of the cameras in the panorama. It has also allowed us to construct high resolution RGBD panoramas where the color and the depth are mutually registered as shown in Fig. 14. In future work, we plan to leverage the relationship between the color and the depth imagery to allow for the extrapolation of missing depth values and an enhanced semantic analysis of the scene.



Fig. 12. A mobile robot equipped with a Kinect One depth camera was used to acquire the raw data used to construct the panoramic depth representations.

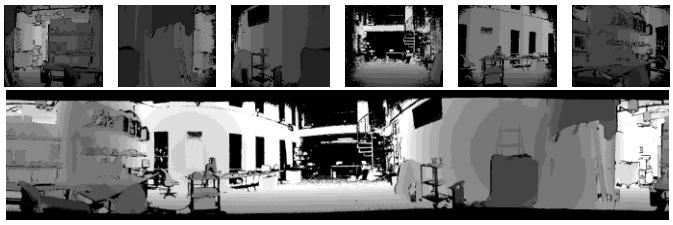


Fig. 13. Depth panorama constructed by fusing a collection of RGB-D frames.

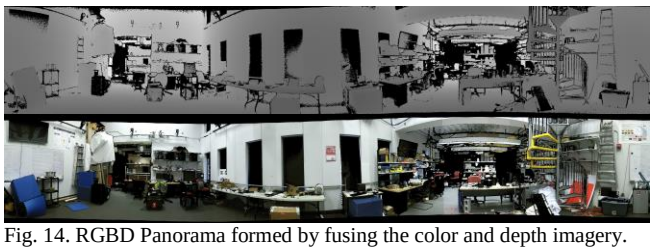


Fig. 14. RGBD Panorama formed by fusing the color and depth imagery. The resolution of the panorama is 8229x1371.

In addition to constructing detailed panoramas of extended scenes we have also made some progress towards developing algorithms that allow for visualization of the captured environment. Fig. 15 shows a screen shot produced by a real-time rendering system which can be used to produce immersive displays of the scene for an Oculus Rift immersive headset.

We have also performed experiments on registering depth panoramas acquired from various vantage points to form a common fused map as shown in Fig. 16. Our goal going forward is to port the panorama construction code we have developed to the embedded computer systems that are currently running on the UAV and UGV so that map

construction can be carried out onboard in real-time. This would also allow us to make use of the onboard sensors to gauge how the sensor moves between vantage points which will help to estimate a better registration between the frames. The objective of this phase is to construct a map consisting of a pose graph where the individual frames are annotated with depth maps. We expect that this representation will provide a convenient basis for reasoning about the 3D structure of the scene. We note that the depth maps provide an efficient implicit representation for the freespace in the scene and we plan to leverage this property to develop efficient algorithms that can produce accurate mesh models of the scene on demand using a variant of the marching cubes procedure.



Fig. 15. Screenshot of real-time visualization system used to explore RGB-D panoramas.

C. Cooperative Localization by Combining Odometry and Ranging Radios

One of the goals of this effort is to be able to map extended environments using teams of MAST relevant UAV and UGV platforms. In order to accomplish that goal we need to consider the problem of estimating the relative position and orientation of robotic platforms that may be operating in disjoint areas. This aspect of the project focused on the problem of establishing the relative position and orientation of a pair of mobile platforms by combining the kinds of measurements that could be derived from odometry or visual odometry system with measurements that could be derived from a pair of ranging radios. This approach is motivated by the observation that these two approaches offer complementary information, visual odometry approaches can be quite successful at establishing the relative trajectory of a platform as it moves through the scene up to an unknown metric scale factor. Ranging radios, such as the recently introduced UWB systems developed by Time Domain Inc. have become smaller and more power efficient and current models and can be used to accurately measure the distance between pairs of platforms. This range information is not sufficient to localize moving platforms in and of itself but when fused with the visual odometry solutions it can establish the metric scale and establish the relative position and orientation of the platforms. Importantly we note that this approach can be used even in situations where the two robotic platforms do not observe any common visual landmarks.

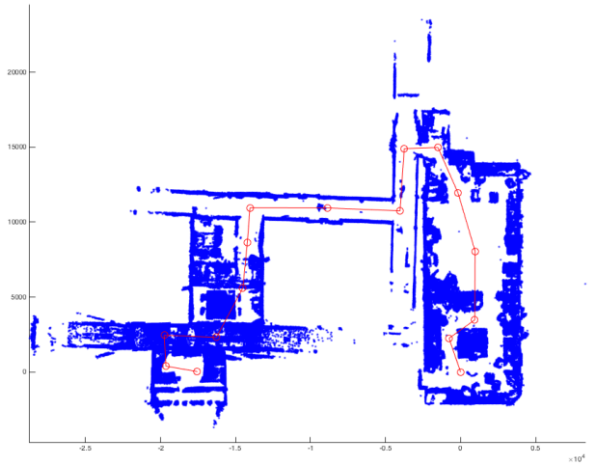


Fig. 16. Overhead view of a map constructed by registering several depth panoramas (red circles) to a single frame of reference.

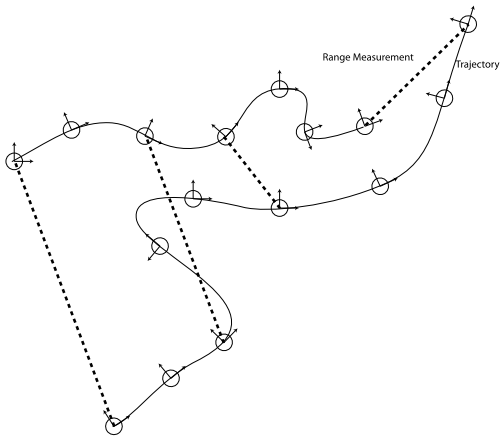


Fig. 17: This figure depicts two mobile platforms which are equipped with odometry or visual odometry systems that allow them to gauge their own trajectories and ranging radios which allow them to gauge their position w.r.t. each other.

Fig.17 depicts the situation we are interested in modeling. In this model we consider a pair of UAVs each executing their own trajectory independently and periodically measuring the range between their positions. Let P_i^1 and P_i^2 denote the positions of platforms 1 and 2 at instant i returned by the visual odometry systems on the platforms. Let λ_1 and λ_2 denote the unknown positive scale factors required to scale the visual odometry solutions to metric units. We assume that each platform is equipped with a 3 axis accelerometer which allows it to gauge its orientation with respect to gravity. That being the case, the frames of reference associated with the two platforms, possibly UAVs, are related by an unknown yaw angle θ and a translation vector **t**. The goal of the relative localization procedure is to recover the unknown parameters, λ_1, λ_2, θ and **t** by leveraging a series of distance measurements between the platforms denote by d_i.

We could imagine phrasing the localization problem as an optimization problem as follows:

$$\min_{\lambda_1,\lambda_2,\theta,\mathbf{t}} \sum_i \{\|\lambda_2 R_z(\theta) P_i^2 + t - \lambda_1 P_i^1\| - d_i\}^2$$

Here the goal is to find the choice of parameters that bring the predicted distance measurements into best agreement with the observed values. The problem with this approach is that the resulting optimization problem is non-convex so one would require an appropriate starting point for the procedure. To this end we observe that the equation relating the predicted distance measurement to the observed distance measurement can be rewritten as follows:

$$\begin{aligned}
\|\lambda_2 R_z(\theta) P_i^2 + \mathbf{t} - \lambda_1 \mathbf{P_i^1}\|^2 &= d_i^2 \\
\left\| \begin{pmatrix} \lambda_2\cos\theta & \lambda_2\sin\theta & 0 \\ -\lambda_2\sin\theta & \lambda_2\cos\theta & 0 \\ 0 & 0 & \lambda_2 \end{pmatrix} P_i^2 + \begin{pmatrix} t_x \\ t_y \\ t_z \end{pmatrix} - \lambda_1 P_i^1 \right\|^2 &= d_i^2 \\
\left\| \begin{pmatrix} -P_{ix}^1 & P_{ix}^2 & P_{iy}^2 & 0 & 1 & 0 & 0 \\ -P_{iy}^1 & -P_{iy}^2 & P_{ix}^2 & 0 & 0 & 1 & 0 \\ -P_{iz}^1 & 0 & 0 & P_{iz}^2 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \lambda_1 \\ \lambda_2\cos\theta \\ \lambda_2\sin\theta \\ \lambda_2 \\ t_x \\ t_y \\ t_z \end{pmatrix} \right\|^2 &= d_i^2 \\
\mathbf{c^T} \begin{pmatrix} \lambda_1^2 \\ \lambda_1\lambda_2\cos\theta \\ \lambda_1\lambda_2\sin\theta \\ \lambda_1\lambda_2 \\ \lambda_1 t_x \\ \lambda_1 t_y \\ \lambda_1 t_z \\ \lambda_2^2 \\ \lambda_2\cos\theta t_x \\ \lambda_2\sin\theta t_y \\ \lambda_2 t_z \\ (t_x^2 + t_y^2 + t_z^2) \end{pmatrix} &= d_i^2
\end{aligned}$$

In the final equation **c** denotes a 12-element vector whose entries are computed from the entries in P_i^1 and P_i^2. This final expression can be derived by carefully considering the structure of the induced quadratic form.

In conclusion, we can view each distance measurement as a linear constraint on a vector whose elements are simple polynomial functions of the parameters of interest. This suggests the following algorithm: collect a set of 12 or more range measurements as the vehicles move through the scene. This yields a linear system that can be readily solved. We can then easily recover estimates for the parameters of interest from the first 7 elements of the recovered vector. The estimates derived in this manner could be used as a starting point for a non-linear optimization process if desired.

We have validated the efficacy of the described algorithm on realistic data in an experiment comparing the relative flight paths of a pair of quadrotors – outputted from our algorithm – with the measurements from a Vicon system. The quadrotors were each outfitted with monocular cameras, IMUs, downward facing laser range finders, and the ranging radios. The robots were then flown in a 3x3 meter Vicon space, facing opposite directions, with their heights kept approximately constant. When compared to the measurements taken by the Vicon system, our algorithm had an average discrepancy of 15.7 cm and a median error of 13.9 cm. The results of this experiment were published and presented in the 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) hosted in Daejeon, Korea. Now that we have properly characterized the localization error in an obstacle-free environment, we will now move towards verifying the system in a more realistic indoor office-like environment.

D. Strategy for Exploration

We extend existing information maximizing exploration strategies for single UGVs to effectively and efficiently traverse the workspace in the context of panoramic RGB-D images. Rather than characterize information content in terms of information per volume in the workspace, we focus on identifying "feature rich" workspace regions identified from the panoramic RGB-D images. Map representations such as occupancy grids provide a probabilistic representation of space that supports exploration methods based on frontiers and map entropy. But while location based mapping is useful for discovering traversable terrain, it is not ideally suited for object identification and situational awareness. Fig. 18 shows two examples of how view angles can significantly impact the information content presented by different angles in the workspace and one's ability to identify objects



(a) (b)

Fig. 18. (a) Figurine of a gnome presents significantly different information when viewed from the back (left) and front (right). (b) A collection of interesting features are difficult to identify at a shallow viewing angle (left) but are easy to spot when viewed head on (right)

In general, the success of identification and retrieval missions requires situational awareness gathered from unstructured environments that may hinder robot operations. Multi-robot control is driven by the needs of high-level tasks that require detailed perception: the individual robot is a means, not an end. As such, our approach leverages the RGB-D images to create lower dimensional representations of the feature space for exploration purposes as well as the navigable free space for navigation purposes. This work combines navigation maps with information maps to guide a robot team to explore and survey an unknown environment. A novel information-spatial representation is used for data acquisition,

and space-efficient RGBD panoramas are used for detailed image and geometry capture.

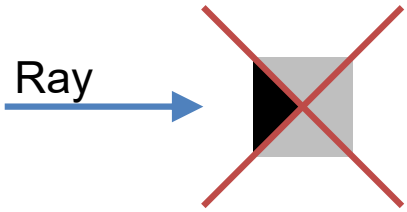


Fig. 19. A sensor beam intercepts a previously unknown cell and the appropriate bin is updated; possible interception angles are organized into 4 bins with the regions being divided by the red lines

The novel representation of space in question is an "angle" occupancy grid: the grid cells record observations of occupancy, but also an inside-out estimate of how thoroughly surface geometry has been observed. An illustration of this concept is shown in Fig. 19. Using the proposed angle occupancy grid to represent the map, we employ a mutual information (MI) exploration strategy.

$$x_t^* = argmax(IM[m; z_t|x_t])$$

where m denotes the map, z_t denotes the measurement at time t, x_t denotes the robot's state/pose at time t, and x_t^* denotes next best exploration pose for the robot. Similar to existing MI exploration strategies, our strategy balances coverage of known space with pushing past frontiers into unknown territory by calculating the expected decrease in uncertainty of the map given a predicted future measurement. Different from existing strategy, our angle occupancy grid computed MI strategy takes into account the view angle of the various features of interest within the workspace. As such, MI identifies potentially valuable locations for sensor measurement.

We employ the Cauchy-Schwarz Quadratic Mutual Information (CSQMI) between a predicted measurement and the map to determine the next best position for the robot to obtain a panorama during its exploration of the space. In general, mutual information can be computationally expensive since it requires integrating the likelihood over entire space, often represented by a grid. Since a closed form expression for the CSQMI exists [21], it enables the quantification of the value of future measurements in a computationally efficient manner. Furthermore, CSQMI has been shown to produce similar results to mutual information [21]. Thus, when coupled with our angle enhanced 2D occupancy grid representation of the workspace for navigation purposes, the result is a scalable exploration strategy with well bounded computational and memory complexities for any given workspace.

1) Simulation Results

The effectiveness of the proposed exploration strategy has been verified in simulation. The images in Fig. 20 show the results of initial simulations. For the environment shown in the left frame a traditional 2D occupancy grid map and an angle occupancy grid map were constructed, middle and right columns. MI was calculated for each free space for visualization purposes and displayed as a heat map. A skeletonization of free space was performed and reachable poses along the skeleton with the highest MI value were identified as targets.

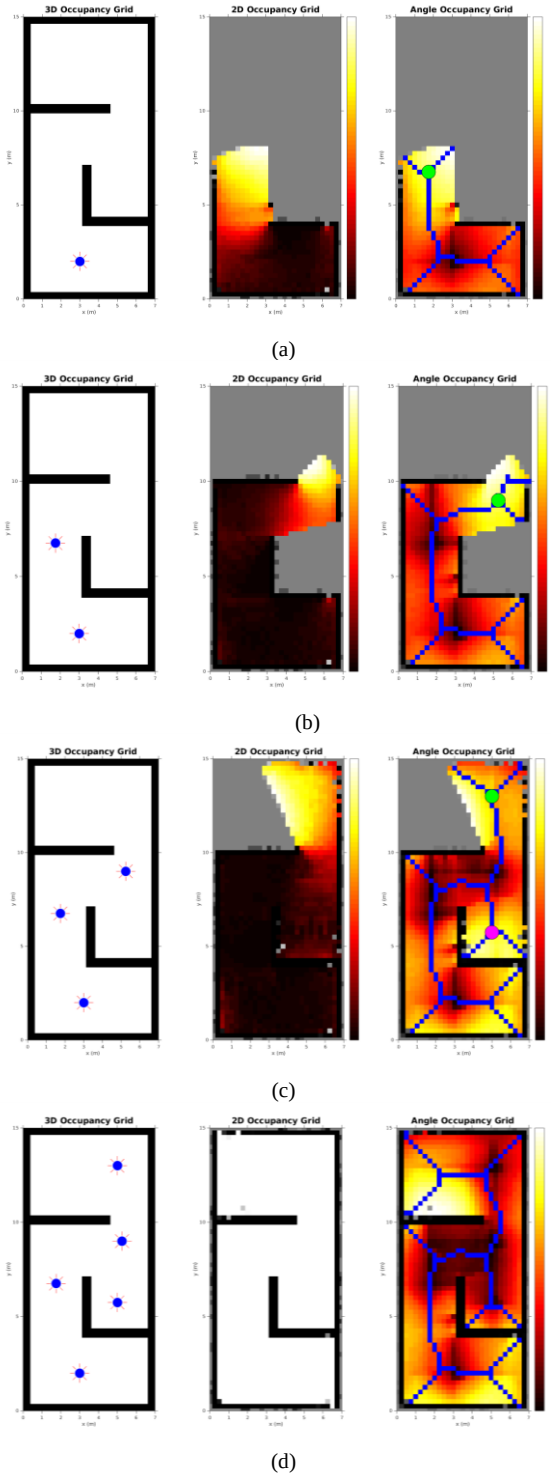


Fig. 20. (a) – (d) Four iterations of the described exploration strategy. In each frame: environment shown with panorama capture locations represented as blue dots (left), 2D occupancy grid with mutual information overlaid as a heat map (middle), and the "angle" occupancy grid with mutual information, skeleton graph (blue), exploration goals (green), and reconstruction goals (pink).

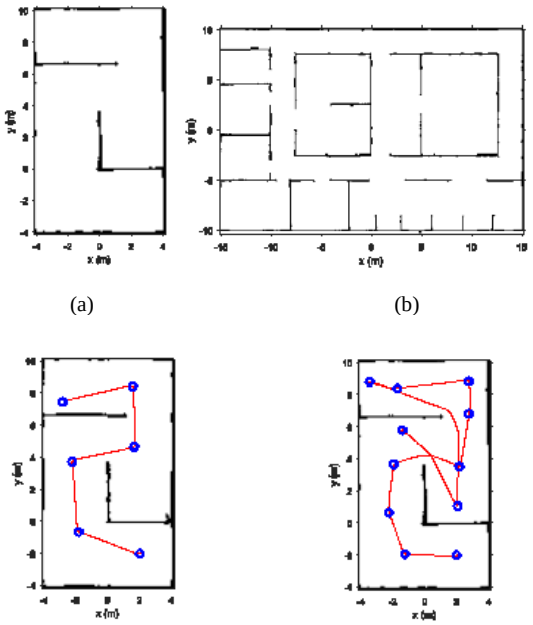
The deployment of multiple robots results in speeding up the exploration process, especially when robots are tasked to focus on distinct regions of the workspace. Our strategy is particularly well suited for small, resource constrained platforms since each robot only has to compute the CSQMI objective function once to determine the next best position to capture a panorama. This enables us develop a simple coordination strategy where the single robot exploration strategy is executed in parallel by a team of robots. Since travel costs are already baked into the exploration strategy, our coordination strategy simply requires individual robots to share their local maps to ensure that goals generated by each robot do not overlap with any past or future panorama capture locations. We note that local maps are only updated each time a panorama is captured and as such they need only be transmitted occasionally, reducing the required bandwidth. The result is a distributed strategy for cooperative mapping of an unknown environment. Critically, this approach relies on the exchange of small occupancy grids rather than accumulated imagery or detailed surface reconstructions. This is the difference between megabytes-per-second and kilobytes-per-hour in terms of communication cost.

In addition, we validated our exploration strategy through extensive simulations in a variety of representative indoor environments with teams of one to five robots and compared our approach with an existing frontier method [22]. We use two simulated environments (see Fig. 21(a) and 21(b)) to compare the resulting exploration locations. An Asus Xtion Pro Live RGBD camera was used to provide observations of the environment. Each panorama comprised of 36 images, one for every 10 degrees of rotation. Throughout our simulations and experiments, we assume the pose of each robot in a global coordinate frame is provided. Since our objective is to validate the proposed exploration strategy, we assume robot localization is solved since it can be achieved via the previously reported on-board localization using ultra wideband (UWB) radios. Fig. 21(c) and 21(d) show the capture locations of panoramas in a simple indoor environment spanning a space of 8x14 m2. The results are summarized in Table II.

TABLE II. SUMMARY OF THE AREA COVERED AND PANORAMAS CAPTURED BY EACH METHOD IN THE HALLWAY (FIGURE 3(A)) AND OFFICE (FIGURE 3(B)) ENVIRONMENTS.

Environment	Method	Area Covered (m2)	Panoramas Captured
Hallway	Angle Aware	110.9	11
Hallway	Frontier	105.9	6
Office	Angle Aware	572.0	140
Office	Frontier	571.3	38

For the set of simulated experiments shown in Fig. 21(c) and 21(d), the frontier approach fully explored the Hallway environment after collecting 6 panoramas, while the angle aware approach selected 11 locations to faithfully capture all surfaces. With these panorama collection locations, we can calculate that 33.5% of wall segments observed during the frontier exploration strategy were captured by fewer than 20 pixels per decimeter in any given panorama. For the angle enhanced occupancy grid approach, only 11.4% of wall segments were observed below the same threshold resolution. If we evenly subdivide the frontier exploration trajectory to result in an equivalent number of panoramas ("augmented frontier"), 17.9% of wall segments are observed below the resolution threshold. This demonstrates that angle sensitivity in the exploration strategy results in substantially fewer "blind spots" along surfaces in the environment as it has more information with which to plan. A naive increase in panorama locations alone, however, does not result in the same coverage gains.



(a) (b)

(c) (d)

Fig. 21. (a) Curving hallway environment used with 1 robot; (b) larger office environment used with 5 robots; and simulation results for the curved hallway for frontier (c) and angle aware (d) planning methods. The red line shows the path of the robot and the blue circles represents panorama capture locations.

The prevalence of low resolution blind spots in the augmented frontier strategy is visible in the overlayed histograms in the lower-right of Fig. 22. In this figure, the green histogram bars of the augmented frontier observations stand out on the left of the chart. Precisely where these blind spots arise may be seen in the top row of Fig. 22 where the light blue wall segments on the right side of the map approximately one third of the way up the image, for example, indicate poor coverage.

2) Experimental Results

We have also conducted live experiments in the space shown in Fig. 22(b) which covers an area of approximately *5 x 5 m^2*. The two differential drive ground robots shown in Fig. 22(a) traversed the environment to collect panoramas, sharing maps and goal locations as described. Each robot was equipped

with an Asus Xtion Pro Live RGBD camera providing synchronized color and depth frames for the panoramas and spoofing a laser rangefinder for construction of the angle occupancy grid used during planning. All software was written in C++ and executed on an Odroid-XU4 single board computer on-board the robot running Linux and the Robotic Operating System (ROS). Localization was provided by an external motion capture system.

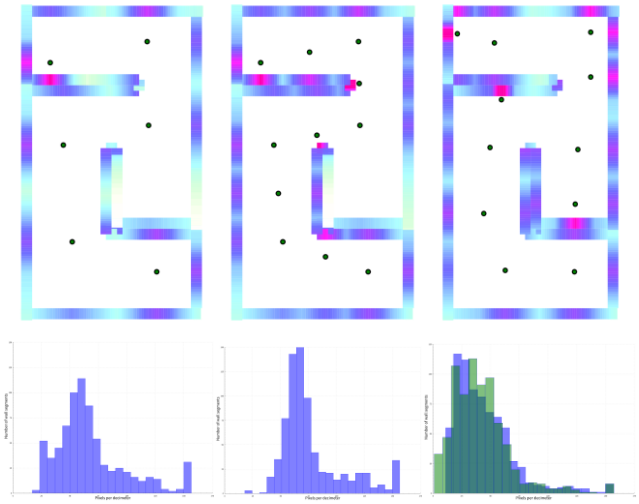


Fig. 22. Top row: map renderings with walls colorized by the resolution at which they were covered for frontier (left), augmented frontier (center), and angle aware (right) approaches. Hue indicates imaging resolution with red indicating high resolution, and blue indicating low resolution. In grayscale, the less saturated wall segments indicate poor coverage. Panorama capture locations are denoted by green circles. Bottom row: histograms of pixels per decimeter for augmented frontier (left), angle aware (center), and augmented frontier overlaid on angle aware (right).

In order to predict future measurements and compute CSQMI, several parameters were specified. Panoramas were modeled as a collection of *360* beams distributed over the interval *[0;2π)*. Given a map resolution of *0:05 m*, increasing the number of beams only resulted in additional dependent beams which intercepted the same cells and would be factored out of the CSQMI calculation. Values of z_{min} *= 0:5 m* and z_{max} *= 4:5 m* were used for the minimum and maximum range of the Xtion, and the noise was modeled with *s = 0:03*.

The results of the experiment can be seen in Fig. 23(c) with the two robots beginning at *(-2;0)* and *(-2;1:5)* respectively. Starting at the same time, the team executed our exploration strategy simultaneously and collected a total of *13* panoramas. Operating in proximity, the team successfully avoided capturing panoramas from the same location demonstrating the effectiveness of our proposed multi-robot coordination strategy.

E. Platform and sensors

Based on algorithms proposed, we want to use suitable hardware that is capable to fulfill the previous mentioned requirements. We propose a quadrotor and hexarotor platform packages. They are equipped with a RGB-D sensor and Google Tango device running onboard vision-based SLAM. In case it is needed, we will design a lightweight, 5 gram carbon fiber shell for both platforms, so that they are resilient to collisions to the wall without causing failure to the propeller blades. The processor unit will be an Odroid platform. See Table 1 for configuration details. Both platforms will be supported by an ARO DURIP grant, awarded to Dr. Kumar and developed in collaboration with Dr. Loianno.

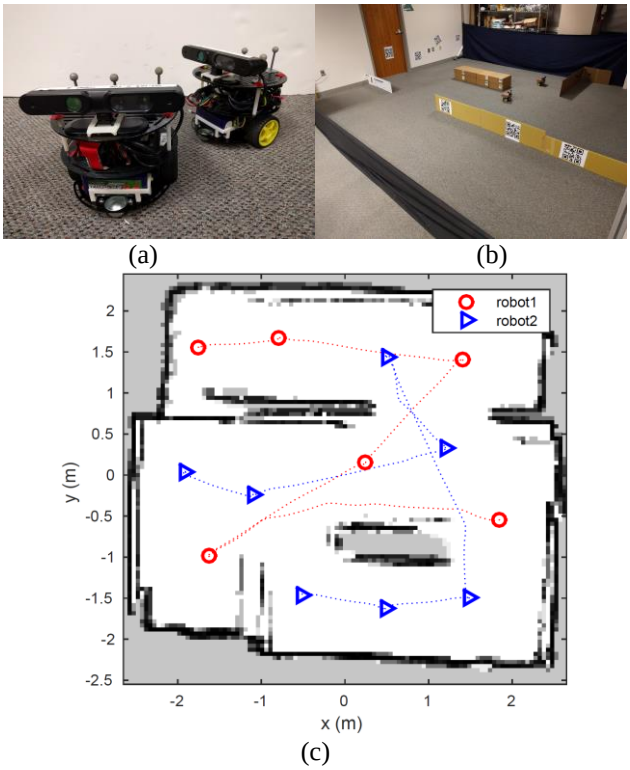


Fig. 23. The ground robots (a), lab space used for live experiments (b), and experimental results showing the panorama capture locations for the two robot team (c).

TABLE 1: CONFIGURATION TO BE EXPLORED FOR VISION BASED LOCALIZATION AND MAPPING

	Package
Processor unit	ODROID-U3, ODROID-XU and ODROID-XUE
Camera	RGB-D sensors and Google Tango
Platform	KMel NanoHex, Asctec Hummingbird and COTS vehicles

V. CONCLUSIONS AND FURTHER WORK

In conclusion, this work addressed questions related to cooperative localization and mapping for autonomous operation in complex and challenging environments. The efforts focused on addressing the key limitations of methods in the state of the art, by using fast visual algorithms for localization, high-resolution dense reconstruction and by employing cooperative strategies with multiple vehicles. The contributions include:

- The development of autonomous micro aerial vehicles equipped with RGB-D sensors and camera phones capable of autonomous, on-board localization, and providing a sparse and dense map of the environment.

The vehicles use their onboard IMU and RGB-D sensor to navigate autonomously.

- The use of depth panoramas in the construction of detailed 3D models of extended environments. We show how a properly equipped robot can acquire these panoramic images by collecting sequences of depth images with a commodity depth sensor. These sequences are stitched into RGBD panoramas that efficiently capture scene information while reducing noise in the captured imagery. Scene structure is then extracted from the RGBD panoramas and used to register a collection of panoramas to obtain coverage of extended areas. The strategy maintains fine geometry detail in 3D reconstructions, uses a modest amount of memory, and enjoys excellent scaling characteristics across large environments.
- A multi-robot information based exploration strategy that maximizes the Cauchy-Schwarz Quadratic Mutual Information (CSQMI) on a novel angle enhanced occupancy grid to guide robots in the collection of RGBD panoramas. The angle enhanced occupancy grid emphasizes perspective in addition to coverage which yields higher quality maps when compared to traditional occupancy grid methods. Compared to frontier exploration, the angle enhanced information maximizing strategy yields better coverage of the space from various view angles.
- An approach to recover the relative position and orientation of a pair of moving platforms by fusing the kinds of estimates provided by visual odometry systems with distance measurements obtained from ranging radio systems. Combining these two complementary sources of information provide estimates that could not be obtained using either source separately. In addition, we show that the localization problem can be posed as a convex optimization problem where the structure of the problem is leveraged to provide an efficient and effective solution. Experimental results demonstrate that the proposed method performs creditably on actual robotic platforms.

The research topics range from cooperative localization and mapping for autonomous operation in complex and challenging environments. In particular, the work focuses on addressing the key limitations of methods in the state of the art, by using fast visual algorithms for localization, high-resolution dense reconstruction, and by employing cooperative strategies with multiple vehicles. The techniques, developed through this thrust, are highly relevant as tools in an autonomous aerial vehicle's repertoire and will be transitioned to the ARL aerial vehicles.

An immediate direction for future work is the comparison of 3D reconstructions of the environment using RGBD panoramas with existing techniques. Specifically, the team has successfully developed algorithms for 3D reconstructions of RGBD panoramas. To quantify the amount of actionable information, the team will experimentally compare the maps synthesized using traditional lidar-based SLAM and from RGBD panoramas. This will be accomplished by running object recognition on the synthesized maps. Known objects will be placed throughout the environment to be explored. A single robot will first explore the environment by executing basic wall following with its visual camera mounted facing the wall. The idea is to guarantee that the robot will always have good views of objects to be recognized. Similarly, the robot will then explore the environment by performing traditional lidar-based SLAM where the visual camera will be mounted facing forward, i.e., in the direction the robot is moving. Lastly, the robot will explore the environment and collect RGBD panoramas using the proposed strategy. We will analyze the number of positive identifications for each case, in situations where objects were incorrectly and/or falsely identified, we will compare the perspective in which the objects were viewed to determine the cause of the incorrect identification. A manuscript on this work is current under preparation.

REFERENCES

[1] S. Weiss, D. Scaramuzza, and R. Siegwart, "Monocular-SLAM-based navigation for autonomous micro helicopters in GPS denied environments", Journal of Field Robotics, vol. 28, no. 6, pp. 854–874, 2011.

[2] R. Voigt, J. Nikolic, C. Huerzeler, S. Weiss, L. Kneip, and R. Siegwart, "Robust Embedded Egomotion Estimation", IEEE International Conference on Intelligent Robots and Systems, San Francisco, USA, 2011.

[3] S. Shen, N. Michael, and V. Kumar, "3D Indoor Exploration with a Computationally Constrained MAV", IEEE International Conference on Robotics and Automation, St. Paul, USA, 2012.

[4] S. Shen, N. Michael, and V. Kumar, "Autonomous Multi-Floor Indoor Navigation with a Computationally Constrained MAV", IEEE International Conference on Robotics and Automation, Shangai, China, 2011.

[5] P. Henry, M. Krainin, E. Herbst, X. Ren, and D. Fox, "RGB-D mapping: Using depth cameras for dense 3D modeling of indoor environments", International Symposium on Experimental Robotics (ISER), New Dehli, India, 2010.

[6] N. Engelhard, F. Endres, J. Hess, J. Sturm, and W. Burgard, "Realtime 3D visual SLAM with a hand-held RGB-D camera", Workshop on 3D Perception in Robotics at the European Robotics Forum, V¨aster˙as, Sweden, 2011.

[7] C. Audras, A. Comport, M. Meilland, and P. Rives, "Real-time dense appearance-based SLAM for RGB-D", Australasian Conference on Robotics and Automation, Melbourne, Australia, 2011.

[8] R. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. Davison, P. Kohli, J. Shotton, S. Hodges, and A. Fitzgibbon, "KinectFusion: Real-time dense surface mapping and tracking", International Symposium on Mixed and Augmented Reality (ISMAR), Basel, Switzerland, 2011.

[9] C. Forster, S. Lynen, L. Kneip and D. Scaramuzza, "Collaborative Monocular SLAM with Multiple Micro Aerial Vehicles", IEEE/RSJ Conference on Intllingent Robots and Systems, Tokyo, Japan, 2013.

[10] D. Fox, W. Burgard, H. Kruppa, and S. Thrun, "A probabilistic approach to collaborative multi-robot localization", Autonomous Robots, 2000.

[11] R. Rocha, J. Dias, and A. Carvalho, "Cooperative multi-robot systems: A study of vision-based 3-d mapping using information theory", Journal of Robotics and Autonomous Systems, Vol. 52, No. 3–4, 2005.

[12] A. Howard, G. Sukhatme, and M. Mataric, "Multi-robot mapping using manifold representations", Proceedings of the IEEE-Special Issue on Multi-robot Systems, 2006.

[13] N. Trawny, S. Roumeliotis, and G. Giannakis, "Cooperative multirobot localization under communication constraints, IEEE International Conference on Robotics and Automation, 2009.

[14] M. Cognetti, P. Stegagno, A. Franchi, G. Oriolo, and H. H. Buelthoff, "3D mutual localization with anonymous bearing measurements", IEEE International Conference on Robotics and Automation, St. Paul, USA, 2012.

[15] Z. Danping and T. Ping, "CoSLAM: Collaborative Visual Slam in Dynamic Environments", IEEE Transactions on Pattern Analalysis Machine Intelligence, vol. 35, No.2, 2013.

[16] J. Solà, A. Morin, and, T .Vidal–Calleja, "Fusing Monocular Information in Multicamera SLAM', IEEE Transactions on Robotics, Vol. 24. No. 5, pp. 958–968, 2008.

[17] T. Lee, M. Leok, and, N. H. McClamroch, "Geometric Tracking Control of a Quadrotor UAV on SE(3) ", IEEE CDC International Conference on Decision and Control, Atlanta, 2010.

[18] T. Lee, M. Leok, and, N. H. McClamroch, "Nonlinear Robust Tracking Control of a Quadrotor UAV on SE(3) ", Asian Journal of Control Vol. 15, No. 2, pp. 391–408, March 2013.

[19] G. Vogiatzis and C. Hernandez, "Video-based, Real-Time Multi View Stereo," Image and Vision Computing, vol. 29, no. 7, 2011.

[20] C. Forster, M. Pizzoli and, D. Scaramuzza, "SVO: Fast Semi-Direct Monocular Visual Odometry", IEEE International Conference on Robotics and Automation (ICRA), Hong Kong, 2014.

[21] Benjamin Charrow, Sikang Liu, Vijay Kumar, and Nathan Michael. "Information-theoretic mapping using Cauchy-Schwarz Quadratic Mutual Information," 2015 IEEE International Conference on Robotics and Automation (ICRA), pages 4791–4798, 2015

[22] Brian Yamauchi. "Frontier-based exploration using multiple robots," in Proceedings of the second international conference on Autonomous agents, pages 47–53. ACM, 1998.

2017 Final Report Metrics

Task Number: SPP-15-2.5

Metrics

Notes: *Please list metrics for the reporting quarter only and only Publications that are accepted and searchable on web.*

Awards and/or Recognitions this Quarter

None

Refereed/Journal Papers accepted this quarter

None

Conference Presentations/proceedings presented this quarter

None

Joint Papers with ARL presented or accepted this quarter – duplication with journals/ conferences expected.

None

Invention Disclosures for this quarter

None

Patent Applications Filed this quarter

None

Patents Awarded this quarter

None

PhDs Award this quarter:

None

Masters Awarded this quarter

None

Students supported fully or partially (by name)

Other: Daniel Mox
MS: *Name(s)*

PhD: Armon Shariati
Post Docs: *Name(s*
Master Degree: *Name(s*
PhDs Awarded: *Name(s*

MAST CTA Member Rotations into ARL

None

ARL Researchers Working at Member Sites

None

Workshops and Seminars conducted:

None

MAST Quarterly Metrics Summary Table (Numbers in this table should correspond to the list above):

MAST Final REPORT METRICS TABLE (The data entered into this table should be for the reporting quarter only)	
Awards (Miscellaneous Recognitions and Awards)	0
Refereed/Journal papers	0
Conference presentations/proceedings	0
Joint Papers with ARL	0
Invention disclosures	0
Patent applications filed	0
Patents awarded	0
PhDs Award	0
Masters Awarded	0
Post Docs Supported	0
Master Degree Students Supported	0
Undergraduate Students Supported	0
CTA Member rotations into ARL	0
Visiting Researchers from ARL to the Member Sites	0
Employees hired by ARL from CTA students supported	0

Energy Efficient Approaches for Autonomous Navigation of MAST Platforms

David Blaauw, Electrical and Computer Engineering,
University of Michigan, Ann Arbor
MI, United States
blaauw@umich.edu

Dennis Sylvester, Electrical and Computer Engineering,
University of Michigan, Ann Arbor
MI, United States
dmcs@umich.edu

***Abstract*—Energy efficient on-the-fly storage and processing are essential to realizing vision based navigation of micro autonomous vehicles. Conventional embedded processing units consume high power, achieve low throughput and are not efficient for MAV applications. We propose to integrate various low-power and variation tolerant circuit techniques to enable high-performance, low power and energy efficient processing and storage on MAV platforms.**

***Index Terms*—MAVs, ASIC, VLSI, low power, energy efficient, SRAM, Digital signal processing, Razor, and Variation tolerant.**

I. Introduction

Energy efficient digital signal processing and storage are key challenges to realize real-time autonomous navigation on battery powered micro aerial vehicles (MAVs). Conventional embedded processors consume high power (~W) and have very low energy efficiency on MAV applications (~100Gop/W). ASICs (application specific integrated circuits) with low power circuit and architecture techniques fit the realm of MAVs because of their small size, light weight and high energy-efficiency. In this task, we investigated various low-power, variation-tolerant circuit design techniques and architecture-circuit co-optimizations to enable energy efficient processing on MAVs.

We investigated various methods for energy-efficient processing. We investigate circuit techniques that can be applied on general purpose processing units and also design special purpose modules to accelerate kernel functions in autonomous navigation. Circuit optimizations for general purpose processing include: (1) Ultra-low Voltage Asynchronous Logic, (2) Low Vdd SRAMs. Algorithm-architecture-circuit co-designed ASICs include: (1) a super-pipelined multiplier; (2) a FFT core; (3) a k-best decoder; (4) a Speed-up-robust-feature (SURF) core; (5) a face recognition core; (6) a depth estimation core.

II. ULTRA-LOW VOLTAGE ASYNCHRONOIS LOGIC

Lowering voltage is one of the most effective approaches in improving energy efficiency. In order to enable digital circuits to operate at lower voltages, bubble razor is developed to eliminate timing margins and guard bands in digital circuits, because systems that operate at lower voltages are more susceptible to noise and variation. By operating at these lower voltages, energy consumption can be dramatically reduced, which is significant in the MAST application.

The Bubble Razor approach involves transforming a digital circuit to use two-phase latch timing and then adding in error detection and stalling circuitry on the latches' clocks. A novel stall propagation algorithm allows the system to survive timing errors without any replay or rollback logic, while at the same time not requiring chip-wide communication to achieve error correction. To demonstrate the scalability and general purpose nature of the Bubble Razor technique, it was implemented on a commercial ARM Cortex-M3 Microprocessor, which was only modified using synthesis and circuit techniques, leaving the microarchitecture and RTL as provided by ARM.

A. Quantify overhead of latch transformation

In order for the technique as a whole to save energy in the MAST system, there must be a low overhead for transforming

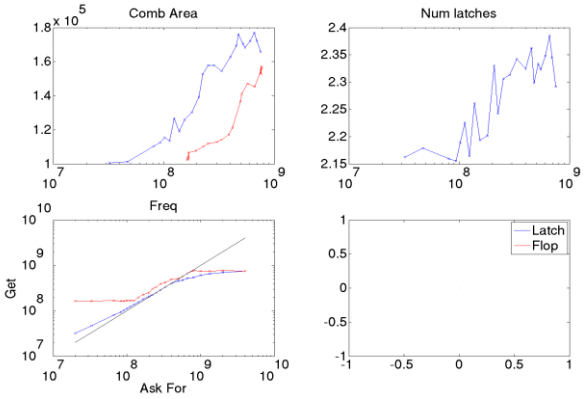


Fig. 1. Latch Transformation Overhead

the circuit to a two-phase latch pipeline. Replacing the flip flops in a circuit with latches and retiming changes the optimization involved with synthesis of the design. We investigated any penalty or benefit of doing this, by applying our technique to the Cortex-M3 Microcontroller, the target platform for our technique. The design compiler was given a range of timing targets to attempt to meet, and some of the results are shown in Fig 1 for both the conventional flip flop based design, and the retimed latch based design.

As is shown in Fig. 1, switching to latches produced between 2.15 and 2.35 latches for every flip flop in the original design. This is close to the ideal number of two latches per flip flop, as a flip flop is made up of two latches. In addition to this overhead, we also observed approximately a 20% increase in circuit area, although as more aggressive timing was achieved, this overhead reduced. There was a negligible effect on circuit performance, as both the flip flop and latch circuits hit their maximum frequency at approximately the same point.

B. Optimize algorithm to assign latches to groups

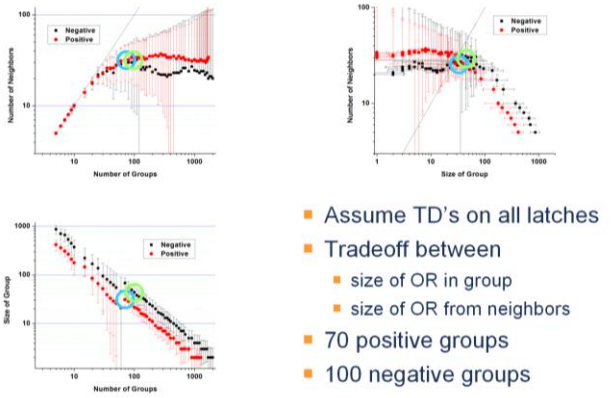


Fig. 2. Results of Group Assignment

The overhead of performing the Razor technique on a latch by latch basis would be substantial, therefore it is necessary to cluster latches into groups that are all stalled together. An ideal cluster would contain latches that tend to share neighbors, which would reduce the total overhead of the system. To perform the clustering, we used hMETIS, a standard tool for partitioning hypergraphs. Since a cluster must only contain latches of the same phase (positive or negative), hMETIS could not directly be applied on a graph representing the latches in the circuit. Instead, the graph was split into two independent graphs, one containing only positive latches and one containing only negative ones. Connections within these graphs were formed if there was a path through a single latch in the main graph, where the weight of the edge was the number of these paths. The two graphs were then partitioned independently with hMETIS to get the cluster assignment for each latch.

With too few clusters, every cluster is a neighbor of every other cluster. With a large amount of clusters, the mean value of neighbors tends to settle, though the standard deviation and maximum continue to increase. To balance the tradeoff between the size of the cluster and number of neighbors, we selected points near where the values are equal.

C. Implement in silicon on commercial ARM Cortex-M3

The pushbutton transformation system was completed, and includes synthesizing and retiming the processor, generating a latch based design, clustering latches, computing neighbors, and inserting the additional razor circuitry. It also applies the boundary conditions for SRAM mentioned above. The transformation system was successfully applied to the Cortex-M3 microcontroller. The Bubble Razor Cortex-M3 is able to run with timing errors and correct for them. It can run at speeds where the original flip flop system and latch based systems fail and provide incorrect results.

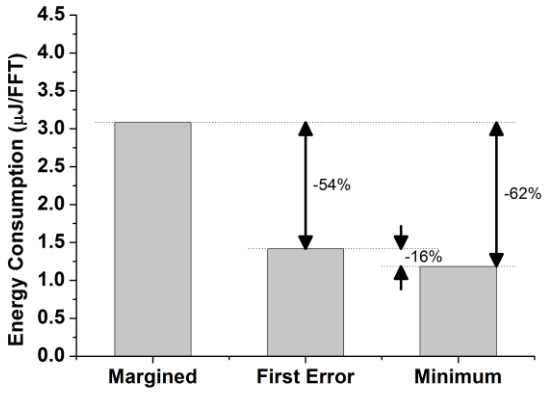


Fig. 3. Performance of Bubble Razor with timing errors present

The chip was further tested to operate correctly under lower voltages where conventional design would fail. For this particular chip, the first timing error occurred at 775 mV. As voltage decreases, the chip operates more slowly allowing more timing paths to fail, causing more detected errors and a higher percentage of lost stall cycles. Initially, as voltage decreases, energy consumption decreases. However, eventually stall cycles become so common that the energy spent on the wasted cycles overtakes the gains from reducing voltage, and total energy consumption increases. In this chip, the minimum energy point occurs at 725 mV. By allowing correct operation beyond the first timing error, energy consumption is reduced by 16%. If the designed was margined

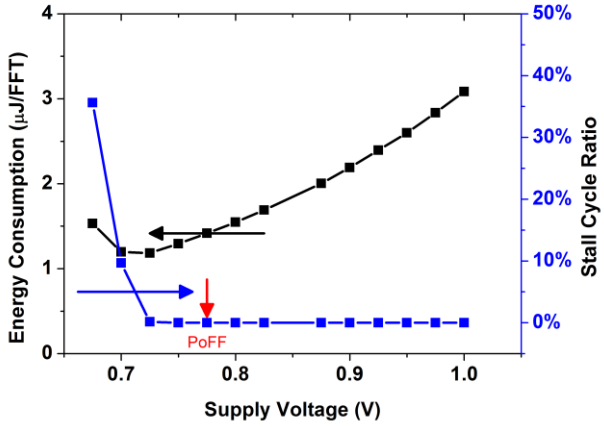


Fig. 4. Performance of Bubble Razor under different supply voltages

such that it operated correctly at 85°C with 10% supply drop, 2σ process variation, and 5% safety margin, the minimum operating voltage is 1V. Therefore, by using this Razor technique, nominal power consumption can be reduced 54% by

operating at the point of first failure, or by 62% by operating at the minimum energy point.

D. Conclusion

We implemented B-Razor on an ARM Cortex-M3 microprocessor without detailed knowledge of its internal architecture to demonstrate its automated nature for MAST GPP platforms. The flip-flop based design was converted to two-phase latch timing using commercial retiming tools; B-Razor was then inserted using automatic scripts. This system is the implemented on a complete, commercial processor, and it provides an energy efficiency improvement of 60% or a throughput gain of 100%.

III. Low VDD SRAMs

Low voltage operation is one of the most effective ways to reduce power consumption due to its quadratic power savings at the expenses of operation robustness and speed. SRAMs are critical digital circuit blocks and they act as bottleneck in low-voltage systems because of their worse voltage scalability compared with logics. Our goal is to develop optimal SRAM bit cells and peripherals for low voltage operation with minimized leakage. We investigate four different SRAMs bit

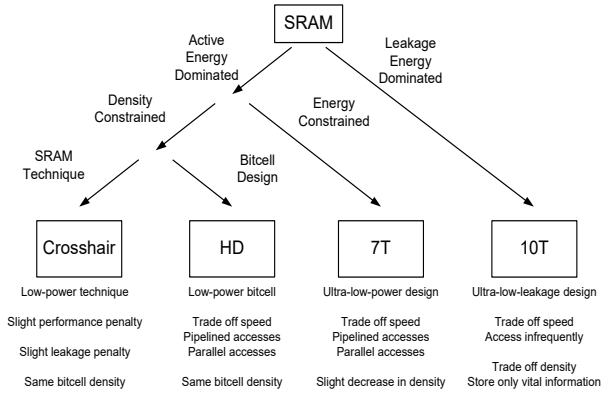


Fig. 5 SRAM Decision Tree

cells shown in Fig. 5 and successfully demonstrate three of them (10T, HD, Crosshair).

A. 10T Low Leakage SRAM

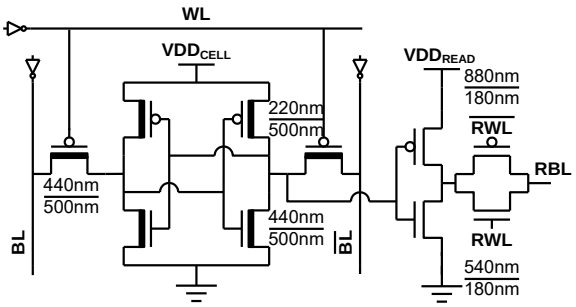


Fig. 6 10T Low Leakage SRAM

In a sensor applications which has a long sleep mode, leakage power consumption in SRAM arrays during sleep mode consumes most of its battery energy. To increase its battery lifetime, it is important to reduce leakage in a SRAM cell. 10T SRAM is introduced for this purpose. Fig. 6 shows the schematic of 10T SRAM. For robust read operation, it has a dedicated read buffer using SVT (standard VTH) devices. However, HVT (high VTH) devices are used in 6T parts which must be turned on even during the sleep mode for data retention. This SRAM was fabricated using 180µm CMOS technology and shows 1.85fW/bit as a leakage power.

B. Cut-off 7T SRAM

To reduce energy consumption more, it is recommended to lower the supply voltage. However, as the supply voltage lowered, SRAM cells are suffer from unstable read and write operation. To enhance robustness in read and write operation, Cut-off 7T SRAM cell is introduced. Cut-off 7T SRAM (Fig. 7) can enhance robustness by cutting of the back-to-back inverter loop.

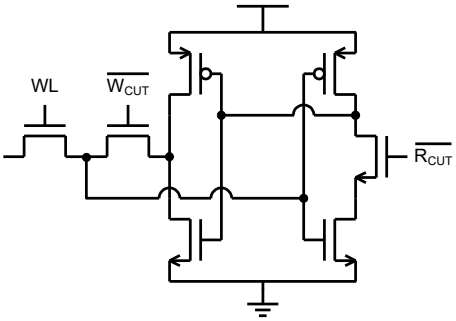


Fig. 7 Cut-off 7T SRAM

C. Half-differential (HD) SRAM

Half-differential (HD) SRAM (Fig. 8) is similar to the traditional 6T SRAM cell. However, it leverages asymmetric

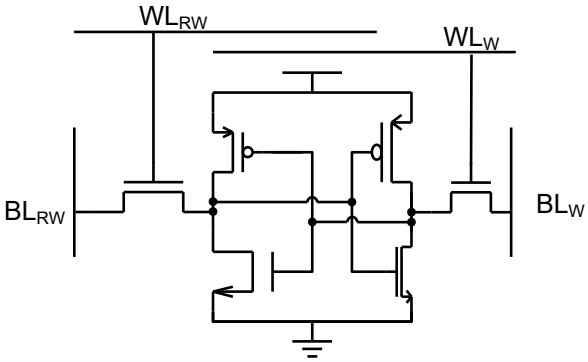


Fig. 8 HD-SRAM

sizing to read through one bit-line and write through two bit-lines. Using this method, we can lower the minimum operating voltage of SRAMs.

D. crosshair SRAM

The crosshair SRAM is shows in Fig 9. The basic structure is similar to the traditional 6T SRAM cell but each bit cell has two independent VDD and GND. Because of process variation, one cell in an array started failing as lowering supplies. To regain robustness of the failing cell, VDD and GND for the cell can be adjusted.

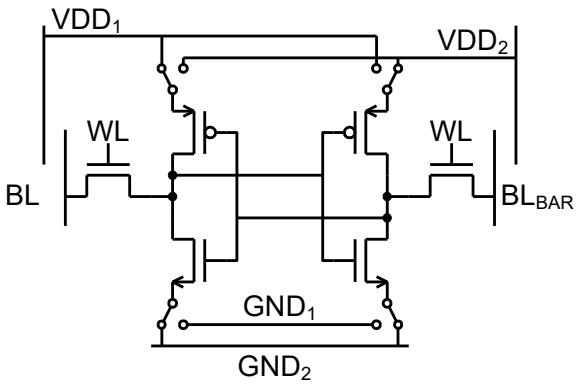


Fig. 9 Crosshair SRAM

IV. ASICs WITH SUBTHRESHOLD DESIGN TECHNIQUES

We investigated architecture-circuit co-optimization method as a low-power circuit technique and our goal is to develop a highly energy-efficient digital signal processing system which is operating in subthreshold region. Such techniques are developed and demonstrated on various kernel functions to demonstrate energy-efficient and variation-tolerance. These include a multiplier, a k-best decoder, a FFT core, a SURF accelerator, a face recognition engine and a depth estimation core.

A. Variation-tolerant Super-pipelined multiplier

The pipeline technique for reduced-leakage operation was applied to the multiplier. We implemented 16bit Baugh-Wooley multiplier which has very regular structure that can be pipelined easily. Although more pipeline stages reduces leakage energy per cycle continuously, the total energy becomes saturated at a specific point due to increasing switching energy of registers. From the simulations, we found that 6-stage is an optimal point for 16bit multiplier. Fig. 10 shows the measurement data of non-pipelined multiplier and proposed multiplier design. The proposed 6-stage pipelined multiplier consumes 30% less energy and shows 1.6X higher performance compared to conventional design.

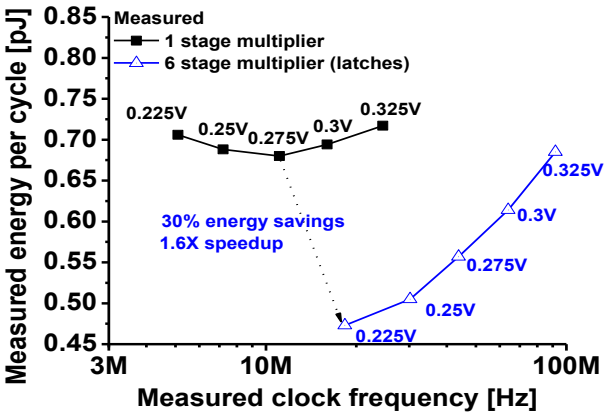


Fig. 10 Performance comparison of the conventional and proposed multiplier design

Although the pipeline scheme saves significant amount of energy, this benefit may be weakened due to wider range of process variation in subthreshold regime. Therefore we also applied latch-based pipeline technique to this multiplier. The latch-based pipelining enables time-borrowing in case of variation and tenders its aspect by spreading increased delay into adjacent stages. Fig. 11 shows the timing failure rate of each design and we can find that the latch multiplier has much higher tolerance against the delay variation and lower cycle delay at the same time.

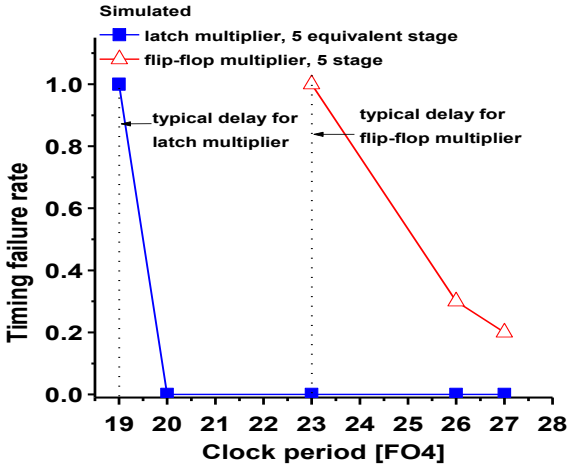


Fig. 11 Multiplier variability

B. Error-resilient K-best decoder

We developed an error-resilient K-best decoder using the timing error model we proposed earlier. This K-best decoder employs the detection-only Razor to detect timing errors and discards the erroneous results. By assigning child node expansion modules adaptively based on the child node survival probability, we achieved the error tolerance up to 30% with only 0.4dB SNR degradation, providing 22.5% energy savings.

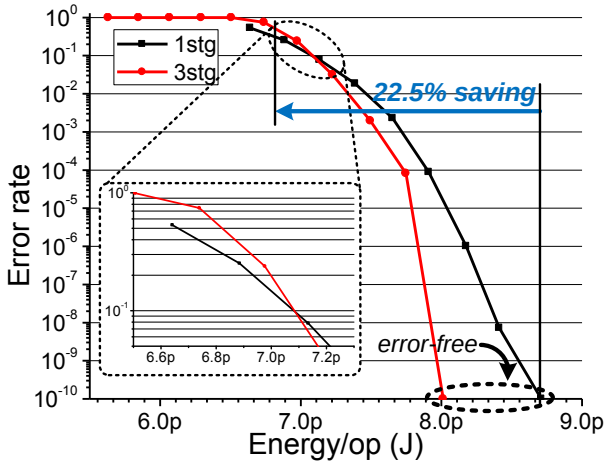


Fig. 12 chip measurements of k-best decoder

C. Energy-efficient high performance FFT architecture

We also designed a FFT core to demonstrate energy-efficient and variation-tolerant subthreshold circuit design techniques. We exploit pipelining method as a low-power circuit technique and architectural modification of conventional pipelined FFT core to enhance energy-efficiency.

The original R4MDC pipelined architecture accepts only one input and waits for all 4 input data required for the first stage radix-4 butterfly to be collected. Each CE is activated only once in 4 cycles period, resulting in CE utilization as low

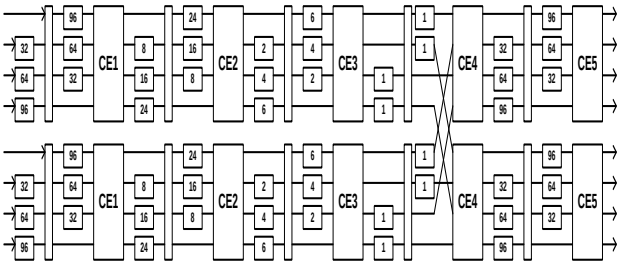


Fig. 13 Modified R4MDC processor (N=1024)

as 25%. However, if 4 inputs are fed at once, we can achieve perfect CE utilization and reduce memory requirements. In addition, we can save more energy by parallelizing the modified R4MDC architecture. We not only increase the throughput but further reduce the number of memory cells per complex multiplier, translating to lower leakage energy. The 2-path modified R4MDC processor is shown in Fig. 13.

This architecture is implemented and chip measurements show that a 1-path modified R4MDC architecture consumes 43.2% less energy with 2.6× better performance than R22SDF architecture at their respective energy-optimal points Vmin. The Vmin decreases from 250mV to 225mV in the 1-path modified R4MDC case due to the reduction of memory leakage. In addition, the 2-path version consumes 27.5% less energy and shows 2× more throughput than the 1-path version with the same Vmin. In sum, the proposed 2-path version has 2.4× better energy efficiency and 5.2× more throughput compared to conventional R22SDF architecture at their energy-optimal points. The results are shown in Figure 14.

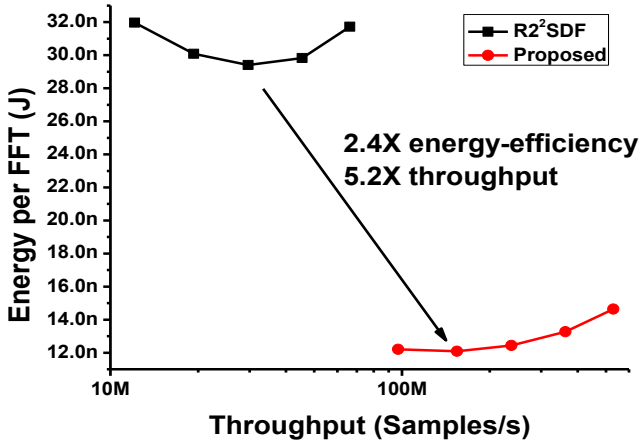


Fig. 14 Energy and throughput of R22SDF and proposed architecture at 200~300mV

D. Energy-efficient high performance Feature extraction

Feature extraction of image is a key step to many visual applications such as object detection, face recognition, surveillance system and visual navigation. It basically detects interest points from given image or video and make a set of feature by describing each interest point. There have been many research activities on this feature extraction, and SIFT (Scale-invariant Feature Transformation) is now widely used

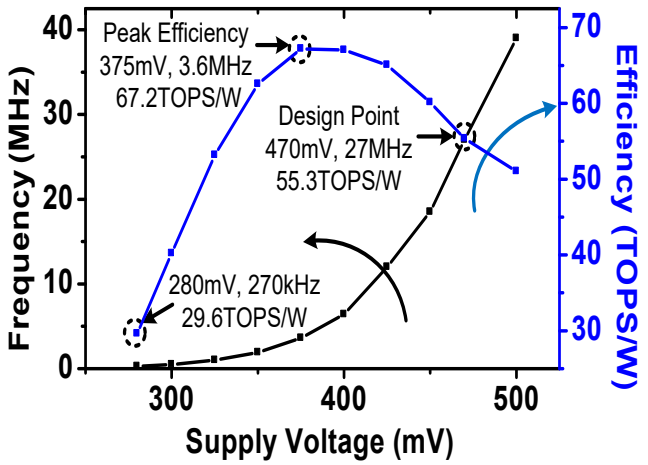


Fig. 15. Performance of the feature extraction processor (N=1024)

due to its robustness and performance. However, well-known feature extraction algorithms including SIFT requires huge amount of computing power and it is nearly impossible to implement them on embedded systems. Although other variants of SIFT such as SURF (Speeded-Up Robust Features) are proposed, even with these algorithms 1GHz embedded core can process ~1 fps.

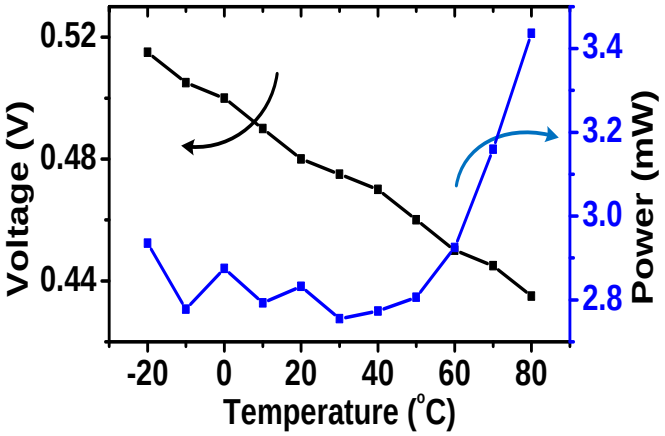


Fig. 16. Measured performance of the accelerator chip with different operating temperatures

We developed a practical hardware solution that can deal with 30fps, 640*480 video, and consuming 2.7mW at 27MHz with a supply voltage of 470mV. The feature extraction core is 3.5x more energy efficient than all prior works. Detailed results and comparisons are shown in Fig. 15 below.

Chip performance under various temperatures are shown in Fig. 16. As temperature increases, accelerator speed reduces and lower operating voltages down to 0.44V become sufficient. However, the power consumption only varies within ~20% meanwhile and it confirms that the proposed accelerator properly operates with the temperature of -20~80°C.

E. Energy-efficient high performance face recognition

Circuit-system co-optimization techniques for energy-efficiency are also demonstrated in the face recognition processor. Because the proposed face recognition system requires a large amount of on-chip memory space (~400kB) and memory blocks occupy most of the area and limits voltage scaling, we propose a novel 5T bit cell structure that has smaller area than conventional 6T bit cell and significantly better read reliability in low operating voltages due to isolated read path.

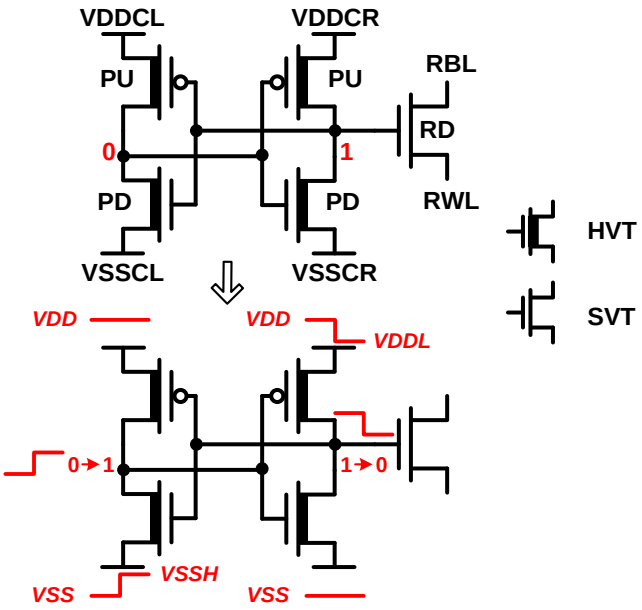


Fig. 16. Proposed 5T bit cell and its write operation

Fig. 16 shows the proposed 5T bit cell and the write operation. It has similar read path through the gate of the access transistor as 7T, but it does not have separate write access transistors. Instead, each bit is reset to 0 or 1 and the write operation occurs in one direction (0 -> 1 or 1 -> 0), which means entire memory block has to be reset before the write operation and write may be done only once. However, most of the memory space is occupied by coefficients of support vector machine and cascade classifiers, and hence they need to be updated infrequently. The overall power consumption is

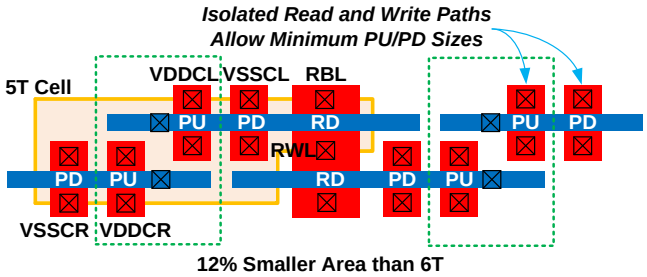
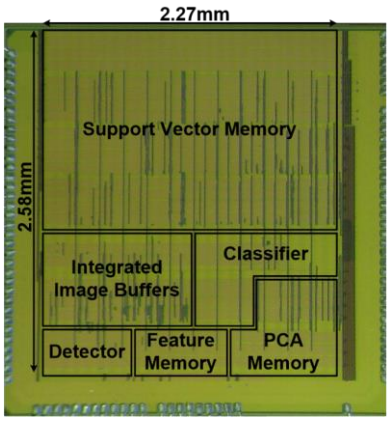


Fig. 17. Layout of the proposed 5T bit cell.



Technology	40 nm
Vdd	600 mV
Clock Freq.	100 MHz
Core Area	2.58×2.27 mm^2
Input Video	1280×720 (HD)
Power	23 mW
Efficiency	4.5 nJ/pixel

Fig. 18. Chip diagram and performance metrics

dominated by continuous read operations during actual face detection and recognition.

Fig. 16 shows the layout of the proposed 5T bit cell. The active region of the read access transistor can be shared with the next bit cell similarly to 7T. It has 12% smaller area than conventional 6T bit cell when it is drawn in a logic rule.

Face recognition accelerator with specialized 5T SRAM for HD (1280×720) images are implemented and tested. The proposed design detects faces from the input image using cascaded classifiers. A SVM (Support Vector Machine) performs face recognition based on features extracted by PCA (Principal Component Analysis). Algorithm optimizations including a hybrid search scheme that reduces the workload for face detection by 12×. A new mostly-read 5T memory reduces bit cell area by 7.2% compared to a conventional 6T bit cell while achieving significantly better read reliability and voltage scalability due to a decoupled read path. The resulting design consumes 23mW in 40nm CMOS (Fig. 18) while processing both face detection and recognition in real time at 5.5 frames/s throughput.

F. Energy-efficient high performance depth estimation core

Precise depth estimation is a key kernel function to realizing autonomous navigation on MAST platforms. State of

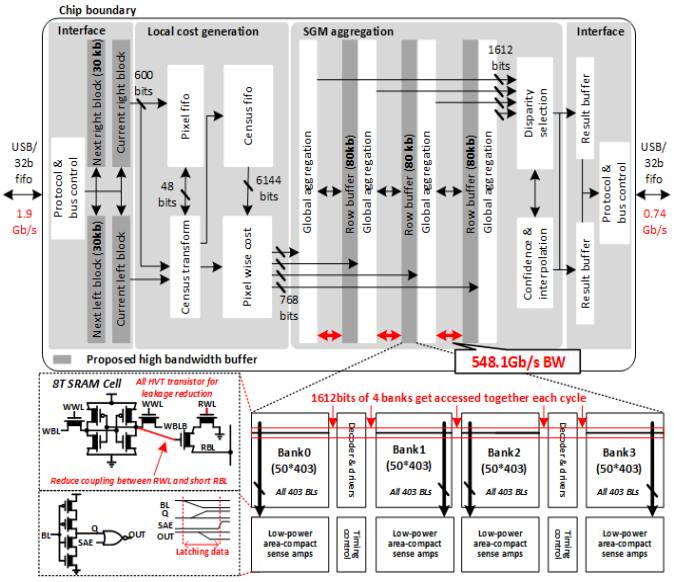


Fig.19. Stereo chip diagram and memory schematic

the art algorithms achieve high accuracy because they effectively handles low texture regions due to their global optimization of the disparity between a left and right image over the entire frame. However, accurate depth estimation involves massively parallel computation (~2TOP/s) and extremely high bandwidth memory access (38.6Tb/s) for 30fps HD resolution. This leads to ~20 sec runtime for an HD image pair on a 3GHz CPU requiring ~386MB memory and >35W power consumption. Together, these factors place it well outside the realm of MAVs. Circuit-system co-optimization techniques are proposed to implement the energy efficient depth estimation core. We uses a new image-scanning stride to enable a deeply pipelined implementation on a block-based Semi-global matching (SGM) processing.

We also designed a customized ultra-high bandwidth strip memory that supports 1.8Tbs access bandwidth. In the proposed design, the row buffers are read and written simultaneously at 170MHz, and all 128 previous aggregated costs are accessed in a single cycle. To reduce leakage power in the 40nm technology, the custom 8T memory bitcell uses HVT transistors. Unlike conventional 8T cells, the read transistor stack is flipped such that the read transistor is not connected to RBL, reducing coupling between RWL and the short, low capacitance RBL. Skewed inverters are used in place of conventional sense amplifiers (1612 per SRAM), reducing sense amp overhead by 2.8×. With the integration of such memory, we are able to parallelize processing and operate at a lower voltage. Architecture diagram and schematic of the high bandwidth memory design are shown in Fig. 19.

The proposed depth estimation processor is fabricated in 40nm GP CMOS. The measurement setup and real-time demonstration platform is mounted on a quadcopter. Real-time image streams captured by the stereo camera are rectified, block-partitioned by the ODRIOD-XU4 board, and then sent to the depth estimation processor through a USB3.0 interface. The proposed processor achieves 7% outlier accuracy on KITTI and an 8× improvement in disparity range compared with prior works. Our system consumes 836mW to process

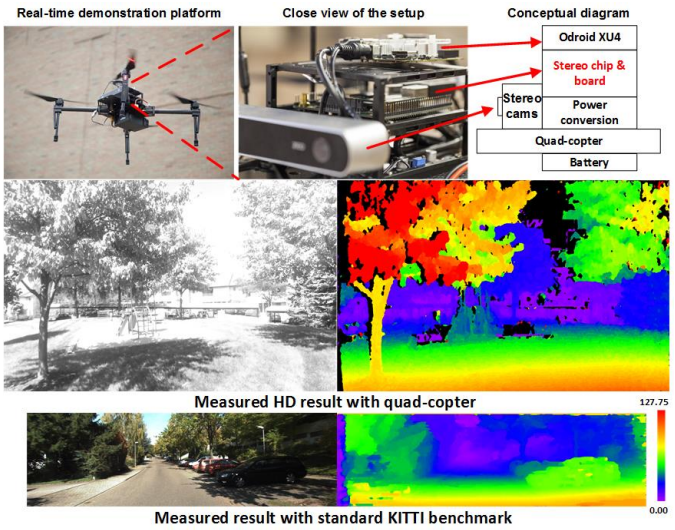


Fig. 20 Stereo chip measurement setup and results

30fps full HD images at 0.0262nJ 'normalized energy', achieving a 5.8× improvement over prior works. Power reduces to 55mW for VGA images at 30fps, yielding 0.0117nJ normalized energy.

G. Energy-efficient high performance 6D vision core

A energy-efficient reconfigurable 6D vision processor is designed for autonomous MAV and VR applications. The processor is designed to achieve 128 levels of depth output at 30fps FHD resolution as well as ±80 levels of optical flow on at 12fps FHD resolutions. Fig. 21 shows the simulated hardware outputs that from the chip.

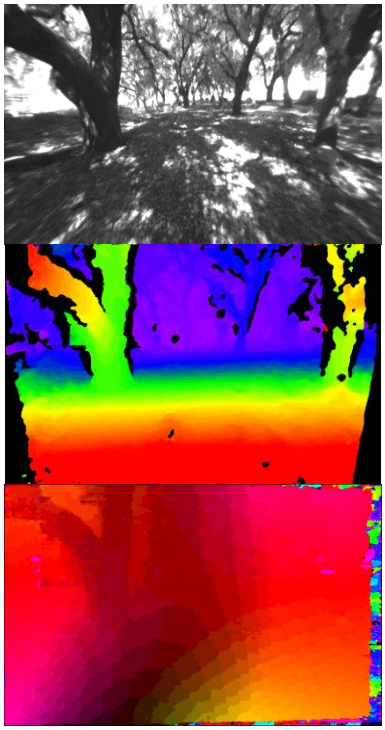


Fig. 21 original frame, depth and optical flow

Proposed hardware architecture is shown in Fig. 22. NG-fSGM_local and NG-fSGM_global are the two major processing blocks. The NG-fSGM_local computes the pixel-wise matching cost, and the NG-fSGM_global module performs cost aggregation. NG-fSGM_local involves fetching all matching candidates from the memory and all matching candidates could have random address. Therefore, it requires the very large memory bandwidth, about 3.6Tbs. Conventional implementation of such large memory access will incur significant latency, which will reduce the frame rate, and consume much power. I observe that in optical flow computation, matching candidates of neighboring pixels coalesce heavily and, on average, only 47 of 138 addresses are unique. We customize the memory access xbar (grey boxes) with the coalescing check to achieve the high-bandwidth as well as coalesce resolution.

Moreover, we partition the processing into two clock/power domains. The red dash line marks the high voltage/frequency domain and the rest of the chip functions at the low

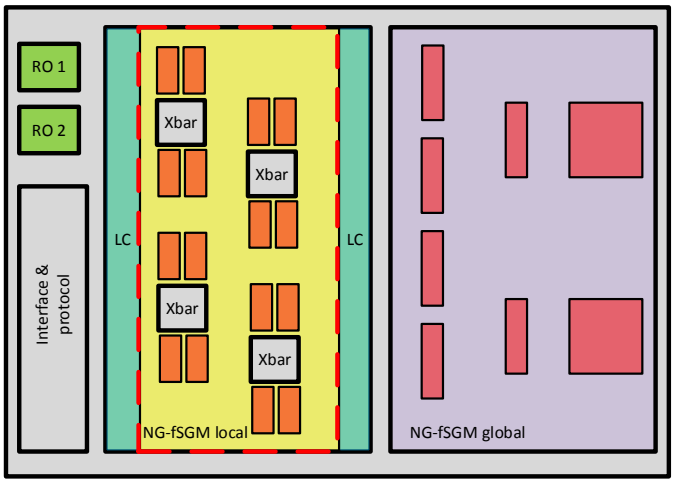


Fig. 22 6D vision chip architecture

voltage/frequency domain. This technique helps improve energy efficiency of the chip.

The Xbar circuit performs arbitration, address transmission and data broadcast. Coalescing check is efficiently resolved at each cross point. In simulation, this coalescing resolving xbar helps improve chip throughput by 79.6%.

This design is taped out in July 2017. We will conduct chip test and system integration late this year.

V. Conclusions and Future Works

During the MAST program we have explored various circuit and VLSI optimization technique to enable energy efficient on-the-fly processing of MAST platforms. We are currently working towards transferring the stereo depth estimation processor and 6D vision processor to ARL for future field test and system integration.

References

[1] D. Fick, G. Kim, A. Wang, D. Blaauw, and D. Sylvester, "Mixed-signal stochastic computation demonstrated in an image sensor with integrated 2D edge detection and noise filtering," *IEEE Custom Integrated Circuits Conference*, 2014.

[2] Mingoo Seok, Dongsuk Jeon, Chaitali Chakrabarti, David Blaauw, Dennis Sylvester, "A 0.27V, 30MHz, 17.7nJ/transform 1024-pt complex FFT core with Super-pipelining", ISSCC 2011.

[3] Dongsuk Jeon, Mingoo Seok, Chaitali Chakrabarti, David Blaauw, Dennis Sylvester, "Energy-optimized high performance FFT processor", ICASSP 2011.

[4] Daeyeon Kim, Gregory Chen, Matthew Fojtik, Mingoo Seok, David Blaauw, Dennis Sylvester, "A 1.85fW/bit Ultra Low Leakage 10T SRAM with Speed Compensation Scheme," IEEE International Symposium on Circuits and Systems (ISCAS), May 2011.

[5] Gregory Chen, Michael Wieckowski, Daeyeon Kim, David Blaauw, Dennis Sylvester, "A Dense 45nm Half-Differential SRAM with Lower Minimum Operating Voltage," IEEE International Symposium on Circuits and Systems (ISCAS), May 2011.

[6] Gregory Chen, Michael Wieckowski, David Blaauw, Dennis Sylvester, "Crosshairs SRAM - An Adaptive Memory for Mitigating Parametric Failures," IEEE European Solid-State Circuits Conference (ESSCIRC), September 2010.

[7] G. Chen, D. Sylvester, D. Blaauw, and T. Mudge, "Yield-driven near-threshold SRAM design," IEEE Transactions on VLSI Systems, pp. 1590-1598, November 2010.

[8] M. Fojtik, D. Fick, Y. Kim, N. Pinckney, D. Harris, D. Blaauw, D. Sylvester, Bubble Razor: An Architecture Independent Approach to Timing Error Detection and Correction, ISSCC, 2012.

[9] Dongsuk Jeon, Qing Dong, Yejoong Kim, Xiaolong Wang, Shuai Chen, Hao Yu, David Blaauw, Dennis Sylvester, "A 23mW Face Recognition Accelerator in 40nm CMOS with Mostly-Read 5T Memory," IEEE Symposium on VLSI Circuits (VLSI-Symp), June 2015.

[10] Dongsuk Jeon, Yejoong Kim, Inhee Lee, Zhengya Zhang, David Blaauw, Dennis Sylvester, "A 470m2.7mW Feature Extraction Accelerator for Micro Autonomous Vehicle Navigation in 28nm CMOS," IEEE International Solid-State Circuits Conference (ISSCC). 2014.

[11] Ziyun Li, Qing Dong, Mehdi Saligane, Benjamin Kempke, Shijia Yang, Zhengya Zhang, Ronald Dreslinski, Dennis Sylvester, David Blaauw, Hun Seok Kim; "A 1920×1080 30fps 2.3TOPS/W Stereo-Depth Processor for Robust Autonomous Navigation", IEEE International Solid-State Circuits Conference (ISSCC), 2017.

Project: SPP-16-17-3.1

Low-Power Readout for Resonant Sensors

Principal Investigators:	Associate Investigators:

Consortium PI: Mina Rais-Zadeh
Affiliation: Univ. of Michigan - EECS Dept.
Address: 1301 Beal Ave, Ann Arbor, MI
48109-2122
email: minar@umich.edu
phone: 734-764-4249
Country of Citizenship: US

Government PI: William D. Nothwang
Affiliation: ARL
Address: 2800 Powder Mill Rd., Adelphi, MD
email: william.nothwang@us.army.mil
phone: 301-394-1136
Country of Citizenship: US

Relevance to Objective Force and ARL

For an autonomous sensing platform, vision is a key function that can be addressed by including a low-power infrared (IR) camera. The IR camera would enable sensing of human objects (for example) day and night and in cases where salt, dust, heavy rain, snow, blowing sand, or petroleum vapor is present; all conditions that would seriously limit optical vision. Other applications of IR detectors related to the objectives of the MAST program include motion detection, proximity sensing test, and detection of and the distance to nearby objects. In our existing MAST project (ending in 2014), we are building the fundamental theoretical modeling, as well as fabrication understanding and expertise in the application of gallium nitride (GaN) based micromechanical resonators as un-cooled (thermal) IR detector arrays. These IR detector arrays can offer noise equivalent delta temperature of better than 5 mK and response time of faster than 30 µsec while consuming significantly lower power compared to available IR detectors. The challenge with GaN-based resonant IR detector specifically and with resonant sensors in general is low-power and fast conversion of frequency to voltage output. Such challenges are mitigated in the new project through developing new readout electronics. This new low-power readout scheme is expected to reduce the total power consumption of the IR detector array by at least 100× compared to existing resonant readout techniques. Although the circuit concept is applied to readout IR sensors, the same technique is applicable to readout other frequency shifting sensors on the MAST platform, such as resonant gyroscopes. By the end of the 3 year program, we will integrate lower-cost and larger format resonant aluminum nitride (AlN) based IR sensor array with the newly developed readout electronics and will demo the complete system. We also plan to have samples ready in Year 3 to send to other MAST performers if needed.

Research Issues Addressed

As part of the SPP thrust and during 2011-2014, our group has developed and demonstrated a new infrared detector using pyroelectric and piezoelectric effect in GaN [1-3]. The developed microresonator IR detectors fall into the category of thermal detectors, the performance of which is basically determined by steady state response, noise relative to the signal produced by ΔT and the time constant of the detector.

$$\Delta T = \frac{P}{G} \qquad \tau_T = \frac{C}{G}$$

Where P is the amount of power absorbed from the source, G is the thermal conductance from the sensing element to a heat sink, ΔT is the temperature rise due to P, τ_T is the time constant and C is the heat capacity of the sensing element. As can be seen, for a thermal IR sensor array, heat capacity, dimensions and fill factor are all important parameters. We demonstrated a parallel array of GaN microresonators coated with an IR absorbing nanocomposite which converts incident IR energy into heat with high efficiency. The rise in the temperature leads to a shift in the frequency response of the microresonators due to combined piezoelectric and pyroelectric effects. IR detection mechanism is based on sensing the shift in the frequency (or the amplitude) of the IR-exposed resonator with respect to a reference resonator which has no IR absorber. The differentially read resonant IR detectors proved to have ultra-small NEDT (lower than 5 mK for size similar to the commercially available uncooled detectors (see Table 1)) but are expected to have high cost. The problem with the cost stems from two factors: 1) GaN growth on silicon is still expensive, despite everyday improvement to the epitaxy technology; 2) GaN resonators cannot be monolithically integrated with CMOS electronics, complicating the readout circuit design and electrical routing. The challenges with GaN-based resonant IR detector specifically and with resonant sensors in general are mitigated in the new project detailed in this whitepaper through developing new readout electronics and by transitioning the technology to aluminum nitride (AlN). The innovative claims are

1. Developing resonant IR detector arrays with 64×64 elements using AlN resonators and integrated with CMOS readout electronics
2. Developing new low-power readout scheme to reduce the total power consumption by at least 100× compared to existing resonant readout techniques. The resonant sensors are usually read out by placing them in an oscillator loop and converting the frequency output to voltage. The problem with such detection scheme is the high power required for oscillation if all resonators in the array are always on. Each oscillator will consume at least 10 μW of power and when using an array for imaging, thousands of these oscillators are operated at the same time increasing the total power only for oscillators to several hundreds of mW. To reduce the power consumption, one can use sequential reading. The problem is that startup time is high specially using high-*Q* resonators, reducing the detection time to several millisecond (in addition to the detection speed of the IR sensor itself). The new readout scheme tackles these issues by using sequential reading that actually makes use of the ringdown time of high-*Q* resonators when triggered with a low-power signal. The same technique can be applied to readout other frequency shifting sensors, such as mass sensors, gas detectors, and resonant gyroscopes.
3. Demonstrating larger format IR detectors that can be used for imaging using AlN microresonators as the sensing element. The advantage of AlN over GaN is that it can be deposited on CMOS electronics directly reducing the burden with signal routing and

lowering the integration cost. AlN sputtering tool is available in house at University of Michigan cleanroom.

4. Technology transition to ARL. ARL has an established process for implementing PZT-based resonators and actuators. Recently, we have proved that such PZT resonators can also be used for IR detection. In the new 3 year program, we will work closely with ARL to tailor the design for better sensitivity. The PZT resonators are monolithically integrated with other PZT actuators on the MAST platform (*e.g.*, on the wing), reducing the total fabrication and integration cost.

Technical Approach

In the next few sections, we explain in detail the proposed low-power readout method of the resonant sensor arrays. The GaN devices demonstrated by our group proved to have orders of magnitude higher sensitivity than other resonant detectors. However, the problem with the resonant detectors has been their power hungry control circuit, which is used to convert the frequency shift to usable voltage output. We propose to develop a new low-power interface IC having a voltage output that makes the IR sensors accessible to other MAST performers for benchmarking and other uses.

The proposed readout circuit is based on exciting the resonator with a signal close its resonant frequency and then counting the ring-down signal that has been produced by the resonator. This method has several advantages over the conventional oscillator-based approach where an amplifier is connected to the resonator to form a loop which forces the sensor to resonate at its resonant frequency. In oscillator-based readout method, there are two conditions to be satisfied to start the oscillations.

- Gain in the feedback loop must be larger than unity.
- Phase shift around the loop must be an integer multiple of 2π.

As a result, the output of the feedback loop is a sustained oscillation with a frequency equal to the resonance frequency. Although this method is simple and widely used, it has some inherent disadvantages. First of all, the output frequency of the oscillator depends on the phase of the oscillator circuitry; hence, the oscillations will not start until the phase shift around the loop is an integer multiple of 2π; this might introduce a high delay, which is often undesired from system perspective. Secondly, the oscillation frequency depends on the phase noise of the amplifier which must be much lower than the resonator for accurate frequency readout. Another problem is the presence of large feedthrough capacitances C_f. With a large C_f, the phase shift of the resonator at the resonant frequency becomes small and poorly defined; as such, the start-up time of the oscillator-based readout is significantly increased. In the proposed method, the resonator is excited with a signal of frequency close to the desired resonance mode for a few microsecond. Then, the excitation is cutoff and the resonator produces an exponentially-damped sinusoidal signal (ring-down signal). From this ring-down signal, the resonance frequency and even the quality factor can be extracted by counting the number of zero crossings before the signal dies out. Figure 1 shows the readout circuit for this method. The readout has two phases, controlled by two complementary pulses PH1 and PH2. When PH1 is high, the resonator is excited by connecting PORT1 to the excitation source and the amplifier is reset together with the rest of the readout circuitry. After excitation, the duration of which determines the initial amplitude of the ring-down signal, PH1 goes low and PH2 goes high, connecting the resonator output (PORT2) to the amplifier. The ring-down signal is amplified and fed to a comparator

which detects the zero crossings of the amplified signal and extracts the resonant frequency. As this is an open-loop readout, there is no need to wait for the oscillations to start; thus, it offers faster operation. The extracted signal can be compared with a reference resonator output to detect the differential frequency shift. This can be done either digitally by saving the results of each measurement and subtracting the reference from them or by an analog method where a mixer is used to down-sample the extracted output.

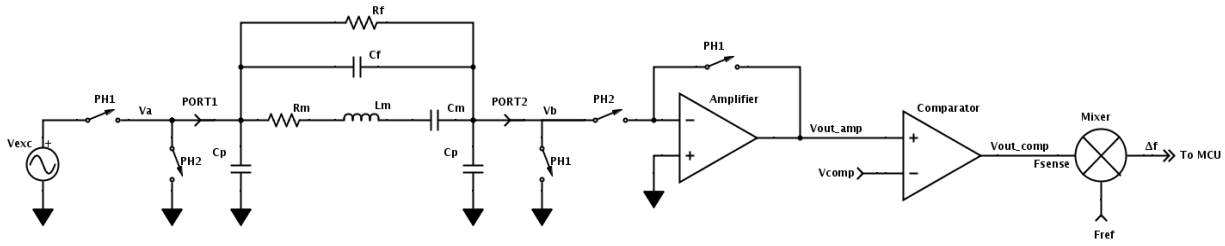


Figure 1. *The proposed low-power readout circuit.*

CMOS Comparator Design: A continuous-time, tunable comparator is designed to count the zero crossings of the ring-down signal. With power consumption of only 60µW, comparator achieves a gain of 58 dB in the frequency band from DC to 70MHz. For the design of the comparator, the circuit model of a 69.9 MHz resonator (shown in Figure 2) was used and the excitation was provided for 2µs with a sinusoidal signal of 100mV$_p$ at 70 MHz, followed by the integration time of 30 µs. The comparator circuit is followed by a "mixer" discussed below.

***CMOS Δf Detector*:** As the comparator output is a pseudo-digital signal, two counters can be used to measure the time for a fixed number of ringdown cycles, *e.g.* 1000 cycles. The 10bits down counter counts a fixed number of ringdown cycles which is reloadable while the high frequency 16bits up counter is fired simultaneously. After the number of cycles has reached 0, a stop signal (**E**nd **O**f **C**ount) will be generated that will stop the up counter from counting. The 16bits output which is directly related to f_{res} will then be latched on the up counter's output.

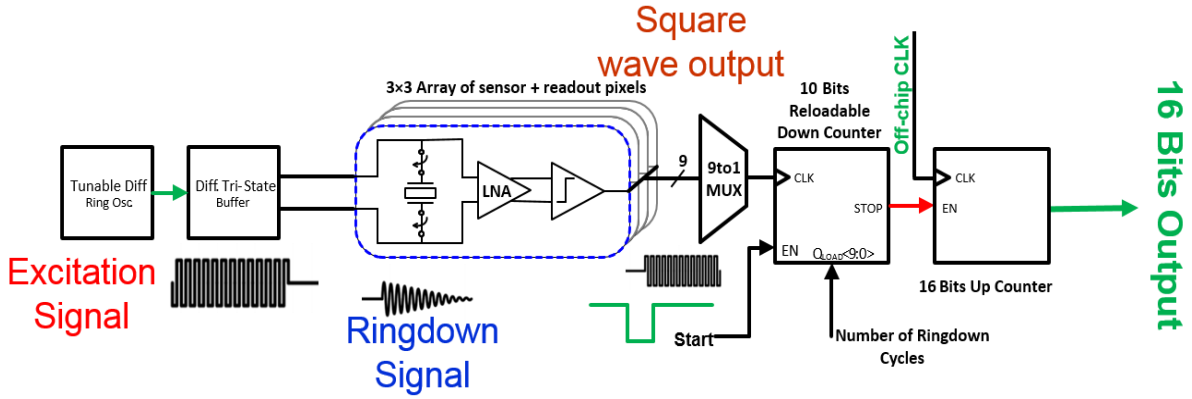


Figure 3. *System architecture (switches and biasing inside the readout channel are not shown for simplicity).*

***System Architecture*:** From system perspective, each resonator in a row will be excited sequentially after being selected by control signals and after the number of ringdown cycles in the down counter has reached 0, the 16bits high frequency counter will stop counting up and the outputs will be latched. The output bits can then be read through the scan chain on the chip. Then, same operation will be performed for the other rows as well. Using such a scheme, the need to have a separate amplifier for each resonator is eliminated. In other words, all sensor outputs can be read by the same channel. To improve speed, multiple channels can also be used and several resonators can be read out at the same time at the cost of increased power consumption. Figure 4 shows the example response of (1,1) pixel in the array using the proposed readout scheme. The measured number of cycles is directly related to the resonance frequency of the selected sensor pixel and thus the change in the resonance frequency of a sensor pixel can be measured by comparing the 16 bits output of each sensor pixel with that of the reference pixel.



Figure 4. *Channel outputs with control signals to read the response of element (1,1) of an array.*

Summary of Project Milestones:

The effort for the resonant detector readout task will include the following three tasks:

1) Develop large format IR resonator array using CMOS-compatible, AlN-based resonators
2) Get the 3×3 array of IR imaging system prototype chip back (expected in early 2018) and do the measurement test and verifications.

Below, the milestones and the performance metrics for the coming year are summarized.

Year 1: In this year of the program, we have designed the readout electronics in CMOS and taped-out the first design. In parallel, we will fabricate large format array of AlN based IR detectors, leveraging our expertise in resonant sensor implementation developed during the existing program. The goal of this year is to demonstrate functional IR detector array of at least 4 × 4 with integrated with the low-power CMOS circuit. The target NEDT is better than 5 mK and target speed is less than 1 msec. The total power for the readout circuit should be less than 5 mW with total detection speed of better than 3 msec.

Year 2: In the second year, we will focus on improving the design of both circuit (for lower power) and resonators (for lower loss). By the end of Year 2, we will have a CMOS readout circuit with measured power less than 1 mW (a 5× improvement over Year 1 goal) and speed lees than 3 msec. The array size maybe limited by the CMOS circuit pad and routing rules. In that case, a separate chip will be fabricated at UM that has larger array of AlN-based IR detectors (of at least 16×16 elements)

Year 3: In this year, we will demo the IR detector array system with the integrated electronics. The main goal of this year of the program is to show a functional uncooled IR sensor integrated in the same package with the readout CMOS circuitry and ready to ship to other MAST performers having power consumption of less than 1 mW. The only element needed along with the developed detector is a lens and the system is ready to be integrated on any MAST platform.

Task Duration: 2015 – 2017
Proposed End State: benchtop sensor with readout electronics (64X64).
End State Demonstration: benchtop comparison of 64X64 sensor with COTS.

References

[1] V. J. Gokhale, J. Roberts, and M. Rais-Zadeh, "High performance bulk mode gallium nitride resonators and filters," in *16th International Solid-State Sensors, Actuators and Microsystems Conference (Transducers)*, Beijing, 2011, pp. 926-929.

[2] V. J. Gokhale and M. Rais-Zadeh, "Uncooled infrared detectors using gallium nitride on silicon micromechanical resonators," *IEEE/ASME Journal of Microelectromechanical Systems*, Dec. 2014.

[3] V. J. Gokhale, Y. Sui, and M. Rais-Zadeh, "Novel uncooled detector based on gallium nitride micromechanical resonators," *Infrared Technology and Applications Xxxviii, Pts 1 and 2*, vol. 8353, p. 835319, 2012.

Recent Alliance Publications, Patent Disclosures, and Patent Applications

Refereed/Journal Papers

1. A. Ansari and M. Rais-Zadeh, *IEEE Transaction on Electron Devices*, in press, Feb, 2014.

2. V. J. Gokhale and M. Rais-Zadeh, "Uncooled infrared detectors using gallium nitride on silicon micromechanical resonators," *IEEE/ASME Journal of Microelectromechanical Systems*, Feb. 2014.
3. V. J. Gokhale, O. Shenderova, G. McGuire, and M. Rais-Zadeh, *JMEMS*, vol.23, no.1, pp.191-197, Feb. 2014.

Conference proceedings *Total 13 including:*

1. M. Moosavifar, A. Ansari and M. Rais-Zadeh, "AlN-on-Si resonant IR sensor array with large temperature coefficient of frequency," *IEEE Sensors Conference*, Oct. 30-Nov. 2, 2016.
2. J.M. Tsai, M. Daneman, B. Boser, D. Horsley, M. Rais-Zadeh, H.Y. Tang, Y. Lu, O. Rozen, A.F. Liu, and F. Assaderaghi, "Versatile CMOS MEMS integrated piezoelectric platform," *International Conference on Solid-State Sensors, Actuators and Microsystems (Transducers'15)*, Anchorage, Alaska, June, 2015.
3. V. Gokhale, C. Figueroa, J. M. L. Tsai, and M. Rais-Zadeh, "Low-noise AlN-on-Si resonant infrared detectors using a commercial foundry MEMS fabrication process," *IEEE MEMS Conference (MEMS'15)*, Jan. 2015.
4. M. Rais-Zadeh, *Proceedings of SPIE: Micro- and Nanotechnology Sensors, Systems, and Applications*, invited, Baltimore, MD, April, 2012.
5. V. J. Gokhale and M. Rais-Zadeh, *Solid-State Sensors, Actuators and Microsystems Workshop*, Hilton Head, SC, June, 2012.
6. S. Yu, V. J. Gokhale, O. A. Shenderova, G. E. McGuire, and M. Rais-Zadeh, *MRS Spring Meeting*, San Francisco, CA, April, 2012.

Patent Applications Filed

None

Patents Awarded

1. *M. Raieszadeh and V. J. Gokhale, "Infrared sensing using pyro/piezo-electric resonators," US Patent pending (US 20120286161 A1), Patent, issued 2015.*

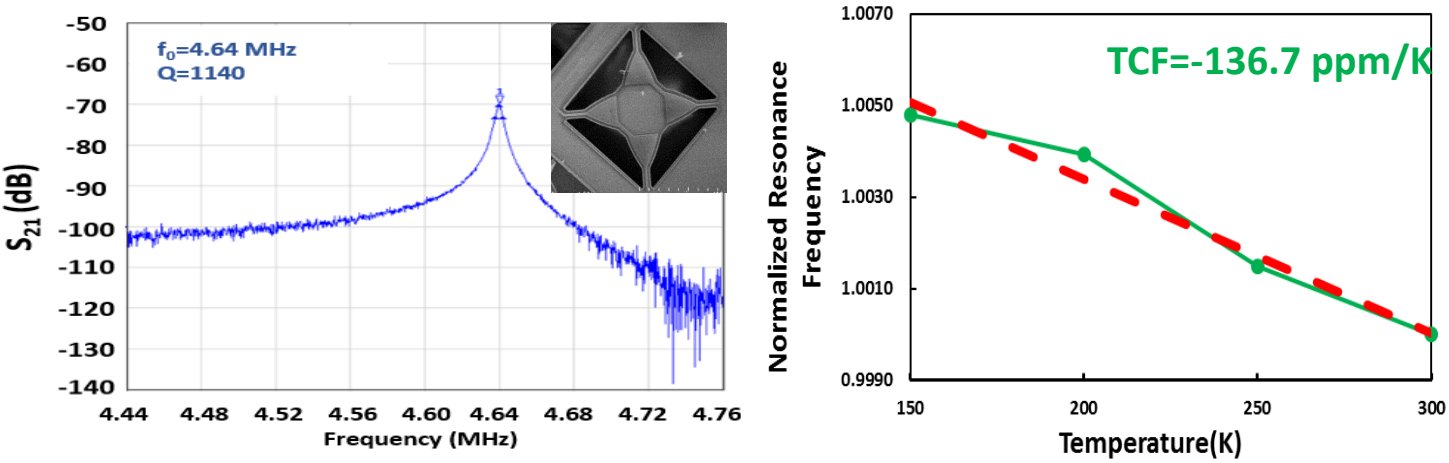
Summary of Collaborations and Support for Cross-Consortia Projects

We have worked with ARL to transition the technology to PZT-type resonant sensors. We have sent our designs to ARL and received fabricated devices from ARL in April 2014. We have tested the samples recently and demonstrated that the PZT resonators could be used for IR detection.

Research Accomplishments:

New Resonator Measurement & Test Results:

New devices with engineered resonator structure which were designed previous year were tested and the TCF (**T**emperature **C**oefficient of **F**requency) of them was measured. These devices were designed to achieve a high TCF in order to improve the IR detection sensitivity. The S_{21} measurement result is shown in figure 5.



(a) *(b)*

Figure 5. *Frequency response of the engineered new resonator resonators (a) Measured TCF of the new devices (b)*

The measurement results showed a large TCF of about -137ppm/K for the new devices which is more than **4×** improved in compare to the similar silicon-based resonators and had not been reported to date.

Readout System Design and Fabrication:

The whole system running for a 3×3 array of sensor pixels has been designed and simulated, with a focus on low power consumption while the imaging system's sensitivity is improved as well. To achieve these goals and improve the immunity to noise, differential signaling has been used on the readout front-end, figure 3. As demonstrated on Figure 4, the 16bits output is generated from each sensor pixel and its resonance frequency is measured using the 16bits digital output. The total power consumption of the designed circuit is 730µW which is below the proposed power dissipation of 1mW. The calculated frequency shift detection accuracy is ~10ppm with the proposed resonator and front end readout circuitry. The CMOS-MEMS chip has been sent out to foundry in June, expected to be back in early 2018. Afterwards the chip's functionality will be tested and verified. Chip layout is shown on figure 6.

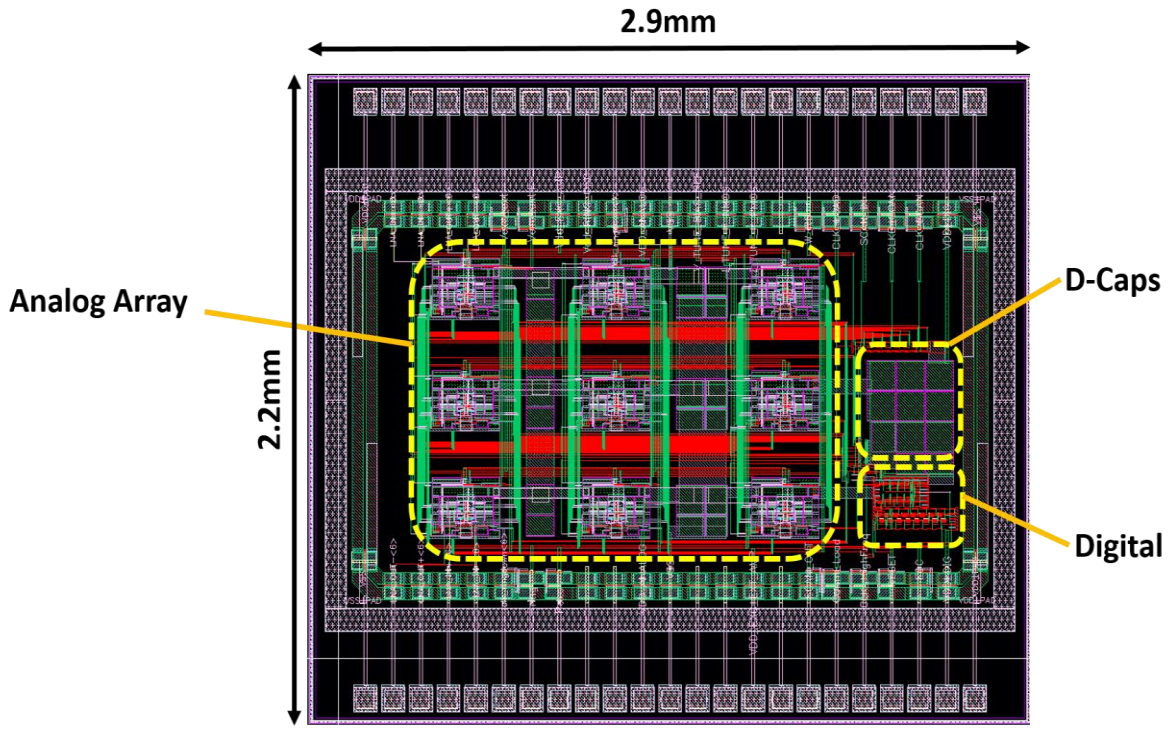


Figure 6. *CMOS chip's top level layout*

The overall performance of the system is summarized in table 1.

Metric	Value
Power Dissipation	730µW
Readout Time (per pixel)	24µs
Min Resolvable Δf	10ppm
Sensor Pixel Size	190µm × 128µm

Table 1. *Readout system's performance summary.*

Summary of Synergy with Existing Work and External Collaborators

Technology Transitions

We are working with InvenSense; as mentioned earlier InvenSense will fabricate the resonant sensor technology using AlN (in their AlN-CMOS shuttle run) that is integrated with GF CMOS 180nm process.

Relation to DoD and Industry Research

We have new funding from ONR and ARO to study the loss mechanism in microresonators.

Metrics

Awards and/or Recognitions this year
Selected as Chair of Hilton Head workshop for 2020 and chair of technical program committee for 2018.

Refereed/Journal Papers accepted this year

None.

Conference Presentations/proceedings presented this year
M. Moosavifar, A. Ansari and M. Rais-Zadeh, "AlN-on-Si resonant IR sensor array with large temperature coefficient of frequency," *IEEE Sensors Conference*, Oct. 30-Nov. 2, 2016.

Joint Papers with ARL presented or accepted this year – duplication with journals/ conferences expected.

None.

Invention Disclosures for this year

None.

Patent Applications Filed this year

None.

Patents Awarded this year

None.

PhDs Award this year

Azadeh Ansari, now assistant professor at Georgia Institute of Technology, Ph.D. awarded in 2016

Masters Awarded this year
Milad Moosavifar

Students supported fully or partially (by name)

MS: *None*
PhD: none.
Post Docs: *None*
Master Degree: *Kaan Tosun, Milad Moosavifar*
PhDs Awarded: *Vikrant Gokhale, Azadeh Ansari*

MAST CTA Member Rotations into ARL – None

ARL Researchers Working at Member Sites

None

Workshops and Seminars conducted:

1. *"GaN for MEMS and MicroSystems," IEEE EDS mini-colloquia, Simon Fraser University, Canada, Oct. 20171.*
2. *"Applications of gallium nitride in MEMS and acoustic microsystems," SiRF, invited, Phoenix, Jan. 2017.*

MAST Quarterly Metrics Summary Table (Numbers in this table should correspond to the list above):

MAST QUARTERLY REPORT METRICS TABLE (The data entered into this table should be for the reporting quarter only)	
Awards (Miscellaneous Recognitions and Awards)	2
Refereed/Journal papers	0
Conference presentations/proceedings	0
Joint Papers with ARL	0
Invention disclosures	0
Patent applications filed	0
Patents awarded	0
PhDs Awarded	0
Masters Awarded	1
Post Docs Supported	0
Master/PhD Degree Students Supported	1
Undergraduate Students Supported	0
CTA Member rotations into ARL	0
Visiting Researchers from ARL to the Member Sites	0
Employees hired by ARL from CTA students supported	0

4. Communication, Networking, and Coordination (CNC) Thrust

Thrust Leader: Vijay Kumar and Alejandro Ribeiro, University of Pennsylvania

ARL Technical Leader: Brian Sadler, Vehicle Technology Directorate

4.1 Research Vision

The CNC thrust concerns the use of communication for carrying out mission objectives, optimizing the network connectivity among the heterogeneous ensemble of robots using the radio, and coordinating the functions of each robot in the ensemble to enable the functioning of the team and optimize the mission objectives. This involves designing highly agile, low-power radios whose performance is modeled online during missions. These models and radio signal measurements can be used for mapping and control to optimize mission objectives, which in turn are functions of connectivity. Connectivity across the team enables individuals to act as a cohesive group enabling coordination at different spatial and temporal scales.

The CNC thrust necessitates advances in radio hardware, models of communication, concurrent design of communication protocols, estimation and control policies, and motion planning to allow robots to act as a cohesive group that is responsive to the operator's commands.

MAST missions will require one or more groups of autonomous and semi-autonomous assets to be coordinated by one or more human operators. MAST groups will need to be responsive to human commands, adapt based on sensed information to provide situational awareness, and be resilient to adversarial conditions. This translates into the following requirements:

1) MAST nodes must be able to establish an ad-hoc communications network between robots, sensors, and Soldiers in indoor and outdoor urban environments, in partially destroyed buildings with rubble, and in caves, with low-power radios and repeaters.

2) MAST must rely on autonomous group behaviors that allow the individual nodes to reconfigure to be responsive, adaptive, and resilient. In particular, the individual platforms must use mobility to enable connectivity, which in turn enables collective behaviors.

3) The individual nodes must be able to alter operational parameters quickly to adapt to changing radio propagation environments, including multipath, congestion, and jamming, as well as adapt to changing throughput and

protocol requirements, and varying range (transmit power) requirements. This can be done by tuning the power, by frequency hopping, altering polarization diversity, adapting routing protocols or controlling mobility. The entire system must be able to optimize its performance by simultaneously changing all these underlying parameters depending on mission specifications.

4) Finally, the communication hardware (radios, antennas, and repeaters) and the processing algorithms must be compatible with MAST size, weight, and payload constraints.

4.2 Army Relevance and Technical Potential

MAST platforms will need to operate with little or no direct human supervision, because groups of platforms will be very difficult, if not impossible, to efficiently manage or control by programming or teleoperation. The CNC thrust has developed mathematical foundations, representations, algorithms and software to enable autonomous operation of the group of MAST platforms.

Autonomy for MAST requires new theory, novel algorithms, and systematic experimentation on many fronts. We must develop a theory and algorithms for inferencing for mobile sensor/robot networks derived from information theory and statistical techniques. We need theory, algorithms, and software for control, perception and learning for mapping 3-D, unstructured, GPS-deprived dynamic environments. We have to create reconfigurable and reprogrammable networks of robots and sensors that will dynamically adapt to disturbances, threats, and redirection and retasking by commanders.

Accordingly, our research program balances basic research on core technologies that are directly relevant to the basic concept of a network of autonomous, adaptive, mobile microsystems with mission-driven technologies representing a unique integration of the core technologies to address MAST scenarios. Our research dovetails closely with the research agendas being pursued by the other MAST thrusts and is informed by the specifics of the concept of operations for MAST.

4.3 Technical Challenges

The vision described in Section 4.1 calls for advances in radio capabilities at the hardware and software level, and algorithms and protocols that can guarantee communication and coordination of the team members. Throughout the execution of the program, we have come to the realization that the team is not a challenge to be managed but an opportunity to be exploited. Consequently, the thrust has been

organized into 3 subthrusts. The first of these thrusts addresses hardware issues in radio design. The second thrust is concerned with the software interfaces for communication and control, whereas the third thrust exploits the communication capabilities of the team to create genuine cooperation in the form of group behaviors.

The primary areas of research and general history of our activities are described for each of these 3 subthrusts.

4.3.1 Novel Radios

Our vision for the CNC thrust calls for the development of a set of radios that operate at diverse frequency bands (20 MHz, 40 MHz, 400 MHz, 2.4 GHz, and 24 GHz). These frequencies were selected based on propagation and technology, and are representative of the different strengths and weaknesses associated with each spectral region, whereas their combined use can ultimately provide a highly flexible and robust system overall. The 20-MHz band has low bandwidth (20 KHz, limited by the antenna) and poor antenna efficiency, but can transmit through many walls, caves, and across terrain. Signals at this frequency do not suffer from multipath fading and can penetrate barriers very effectively. The 400-MHz band can have a 2- to 4-MHz bandwidth with miniaturized antennas developed in the MAST program, and can also transmit through multiple walls with attenuation between 1–3 dB per wall depending on wall type, thickness, and wetness. In this band, there are multipath and some resultant fading. This band is good for indoor to outdoor (and vice versa) direct communication. The 2.4-GHz system has a 20- to 80-MHz bandwidth, but can typically only operate effectively in short-range conditions with relatively less penetration capability. While the antennas and transceivers can be made very small, these radios easily suffer from multipath fading and high attenuation through walls.

Ultimately, it is desired that the radios can be configured to be off until needed, using a wake-up circuit to turn on when a signal is transmitted to them or relying on timing or some other queue. Further, the radios should have a separate transmit amplifier so that the transmitted power can be controlled by the robot to optimize mission parameters while ensuring reliable comms connectivity.

We investigated schemes to increase diversity in frequency, space, polarization, and radiation pattern, enhance data rates, and allow for more robust control of network connectivity. Polarization and radiation diversity will be implemented with proper antenna design to minimize space on small platforms. It is fairly easy to implement diversity, and our research aims for that, whereas commercial radios aren’t necessarily equipped to do so. The transmit frequency, protocol, modulation,

and transmit power can all be configured quickly and easily (1 μs). Further, anti-jam resistance is enabled by variable protocols, modulations, and frequency-hopping schemes.

4.3.2 Communication and Control Software

Our vision for the communication and control software required to interface with the communication hardware (Section 1.1) and the robot platforms (described in the MCE thrust) involves the concurrent design of discrete protocols and continuous control policies to maximize system performance. The main objective is to provide transparent and robust connectivity in support of the mission of the MAST team and co-design control and network management in an optimal and reconfigurable way.

There are many ways to tune and optimize the network characteristics. First, we showed that the positioning of MAST nodes, for example, as relays, can extend the reach of the network and provide sensory information from remote locations via multiple hops. This suggests that mobility of individual nodes can be controlled to maximize performance. In the same vein, we investigated the optimal positioning of repeaters to enhance system performance.

Second, it is possible to integrate the allocation of power and frequency along with the control of mobility and design of routing policies. To do this, we must extend our understanding of radio propagation and build channel state estimation for 3-D environments.

Third, we must exploit diversity along different axes, spatially, temporally, and in terms of frequency. We may also have heterogeneity in terms of small-rate/long-distance links along with large-rate/short-distance links. Again, a system-wide optimization is necessary.

Finally, a systematic approach for exploring and optimizing over the configuration space is necessary. In a d-dimensional world with n nodes, the configuration space is nd-dimensional. Finding configurations that ensure connectivity is hard. Optimization to maximize connectivity is even harder. Clearly, this is somewhat counterintuitive, since a bigger team of robots gives us more degrees of freedom and finding configurations that satisfy connectivity constraints should be easier. It is unclear what the mobility-communications power tradeoff will be on real MAST platforms, but extending (or satisfying) mission lifetimes will definitely be important.

It is important to note that mapping and prediction of wireless channels as a function of positions of individual nodes in the environment is already quite

complicated, especially for the wavelengths that have significant scale fading. This is even more so as we consider the extra dimension of frequency (and power). As we vary the frequency used for communication, the importance of geometric information to the connectivity control problem also varies. At one extreme (24 GHz), visibility-based methods are correct while at the other end (20 MHz), distance-based reactive methods are likely to be adequate. Somewhere in the middle, it is necessary to combine probabilistic modeling and randomized configuration search to find the best positions for intermediate nodes. This requires a unified approach for all these methodologies, with a clear articulation of costs and metrics.

Finally, it is also necessary to employ techniques that allow individual members of the team to take autonomous decisions in environments with intermittent communication. For example, these autonomous actions take the form of decisions regarding which area of the space to explore and when the information acquired is of sufficient interest so as to deserve being communicated back to the command center or other actor. This in turn allows us to endow agents with the ability to reason in the face of uncertainty about what are actions they should take given the partial knowledge they have about the environment, the team, and the command center.

4.3.3 Group Behaviors

MAST requires control technologies and group behaviors for micro aerial and ground platforms focusing on applications related to scenarios discussed earlier. MAST applications will need innovative navigation, control and adaptation strategies to 1) maximize the potential for acquiring data and drawing inferences without direct human control of individual platforms, 2) reconfigure in response to changing environments conditions and commands from the human operator, 3) cooperate and accommodate failures and network breaches, and 4) achieve adaptation and fault-tolerance for safe and reliable operation. Vehicles must be able to localize and map the environment while providing coverage and be able to automatically perform scouting, reconnaissance, and surveillance tasks.

Broadly speaking, group behaviors will be needed for at least these five different tasks:

1) Coverage – deploying platforms and sensors to ensure blanket coverage of a specified area.

2) Robotic routing – realizing a mobile, multi-hop network where a group of robots extend the range of network by automatically adopting positions that will enable relaying of information across the group.

3) Exploration and mapping – enabling a team of multiple robots to automatically explore and map an unknown environment.

4) Tracking – coordinated control of a team of robots to detect, pursue, and track a threat or any object of interest.

5) Persistent surveillance – deploying platforms and sensors to maintain surveillance over a specified area with guaranteed refresh rates.

It is necessary for the team to manage energy resources and adopt strategies that are power aware. While individual nodes may not have the storage capacity to complete tasks, a team should be able to manage resources to accomplish the task with all the constraints. A particularly important focus is on mission op tempo. We are interested in such missions as indoor reconnaissance and search/rescue with time constraints. For example, what are some approaches to deploying a group of vehicles to quickly accomplish a mission in complex 3-D environments that might include the risk of losing platforms while maintaining the desired timeline?

4.4 Technical Plan

The technical plan was executed along the three subthrusts that are described in the subsequent subsections.

4.4.1 Novel Radios

Over the course of the MAST program, we were engaged in development of low-power radio and repeater system architectures for different applications and scenarios. These included high-data-rate radios with beamsteering capabilities at 24 GHz for transferring large amount of data over a short period of time. Radios with very small SWaP and corresponding radio repeaters and very small form factor antennas with pure vertical polarizations were also developed at 2.4 GHz for short-range communication in an indoor environment. The 2.4-GHz system has a 2- to 5-MHz instantaneous bandwidth over an overall 80-MHz bandwidth, but can typically only operate effectively in LOS conditions. Microwave radio repeaters were developed to allow for NLOS operation. While the antennas and transceivers can be made very small, these radios may suffer from significant multipath fading, and high attenuation propagating through walls. To alleviate the path-loss issues in complex environments, the consortium also investigated operation of 400-MHz radios for MAST platforms. The 400-MHz band can have a 2- to 4-MHz bandwidth and can also transmit through multiple walls (attenuation between 1–3 dB per wall depending on wall type, thickness, and wetness). This band is good for indoor to

outdoor (and vice versa) direct communication, although there is multipath and still significant signal fading due to large scatters present in the environment.

More recently, we have focused on communications at much lower frequencies, namely, the lower VHF band centered around 40 MHz. Operation in the 40-MHz band (as well as the 20 MHz band) for MAST has been enabled through invention of super-miniaturized antennas. These small form factor antennas enable relatively small radio packages that can be used on MAST scale platforms. The main advantage of the low-VHF band is very low path-loss and the absence of fast fading that makes communication in complex environments very reliable with low power transmission. This band can be exploited not only for communications, but also for source tracking and geo-location in GPS-denied environment. Highlights of our work include the miniaturized antennas, measurements and propagation modeling in a variety of environments, as well as integration with COTS hardware (based on ZigBee commercial radio) for control experiments. The miniaturized antennas are being enhanced by using non-Foster impedance matching to more than double the antenna bandwidth in the same physical form factor.

4.4.2 Communication and Control Software

The vision for the communication and control software required to interface with the communication hardware (Section 4.1) and the robot platforms (described in the MCE thrust) involves the concurrent design of discrete protocols and continuous policies to maximize system performance. The main objective is to provide transparent and robust connectivity in support of the mission of the MAST team.

To solve these problems, robots must able to build models of communication and networking on the fly and use abstractions that allow efficient routing of messages and maximize the quality of service of the network. This is challenging because RF propagation is notoriously hard to model and requires statistical approaches. However, robots have the potential to continuously learn about their environment and improve their models over time.

Robots must also use their learned models to maximize performance. In particular, they can use abstractions of communication links and their dependency on positions of robots to determine effective placement of individual nodes. By placing the robots optimally and continuously changing the robots' positions to adapt to changing communication links, we can create a cohesive team of robots that are "connected" in the best possible way. Further, it is also possible to find the best power levels and routing policies to maximize performance.

4.4.3 Group Behaviors

The MAST program has paid dividends in many aspects of multi-robot systems. We have developed algorithms for decentralized coverage control in simple 3-D and complex 2-D environments. We have approaches to verify coverage with minimal communication in 2-D environments. We have algorithms and protocols that ensure multi-robot connectivity based on approximate models of communication. Finally, we have developed algorithms and software that enable the planning/control/coordination of vehicles with constraints on actuation and power.

It is clear that we must make additional progress in the area of group behaviors. We must create distributed estimation algorithms and integrate them with control and planning to create a reconfigurable network of MAST robots. We need to enable robots to establish and maintain coverage while subject to communication constraints. This will also lead to capabilities for persistent surveillance. Here, the limitations of available energy storage devices are important to recognize. While our research does not explicitly address the development of novel energy sources, we are actively exploring the best of existing technologies. Further, we must develop energy management and power aware solutions for group mobility that are independent of advances in storage technology. Finally, as the Micromechanics center creates novel platforms with unprecedented agility, we must consider planning and control algorithms for multi-robot systems with dynamics.

This subthrust of the MAST CNC thrust addresses the science, technology, and software infrastructure for group autonomy. In particular, we are interested in the software architecture for individual nodes and interaction across nodes. How human operators will interact with MAST platforms is an important question, one that is relevant for this thrust. We must also consider scaling up to larger teams, incorporating aspects of heterogeneity in sensing, computation, and mobility. Finally, the architecture must support compositionality, task allocation, and deployment, and organization and reorganization into groups or subgroups.

Our work to date has resulted in software architecture for sensing, estimation, mapping, control, and planning for small robots. We have leveraged open-source development and simulation environments for deploying multiple robots. We have made some advances on how knowledge and models can be shared across MAST platforms.

We have further developed systematic approaches to the specification of MAST tasks that can then lead to the automated generation of behaviors for individual and groups of robots. We exploit the potential synergy between aerial and ground

vehicles and realize a cohesive, heterogeneous team. Finally, we have created a set of platforms that function equally well indoors and outdoors.

4.5 Key Program Technical Accomplishments

Over the past decade, the CNC thrust has advanced capabilities that have made it possible to deploy micro-autonomous systems in complex unstructured environments. Since a complete list would be onerous, we report here a few highlights of key technical accomplishments:

1) *Low frequency antenna design:* We have developed high-efficiency, high-bandwidth miniaturized antennas at 40 MHz. These antennas make use of polarization and provide a four-times improvement in bandwidth using non-Foster impedance matching.

2) *Low frequency multi-node networking and geolocation:* We have advanced multi-node networking capabilities using compact low-power, low-VHF radios. These multi-node networks further provide accurate geolocation estimation using low-VHF fusion with other sensing modalities.

3) *Control to enable connectivity and coverage*. We have used models of radio propagation to design motion plans that guarantee connectivity between members of the team. These plans are designed by constructing virtual obstacles in configuration space that represent configurations for which communication is infeasible.

4) *Operation in environments with intermittent communications.* We have studied information games for the design of optimal response systems in environments with restricted communication. The methods rely on evaluating estimates of information value acquired by members of the team.

5) *Optimizing heterogeneous robot networks to maximize quality of service (QoS).* We have exploited heterogeneity in wireless aerial and ground robotic networks to maximize network performance. The heterogeneous composition of the team is leveraged to provide redundancy in space and radio communication use. Methods were developed to exploit synergies between range and bandwidth performance of different communication modalities such as optical and low and high frequency radio.

6) *Communicating information associated with 3-D maps*. The acquisition of maps interplays with the ability to communicate the map to system operators in a time critical manner. We have studied how to construct dense

maps from lightweight sensor platforms for navigation in 3-D environments.

7) *Planning tasks and policies for teams of multiple sensors for situational awareness.* We have produced a framework and specific algorithms for quickly deploying teams of aerial robots in 3-D environments. This accomplishment rests on the solution of problems in dynamic goal assignment and trajectory generation for heterogeneous teams.

8) *Cooperative teams of MAVs exploring complex 3-D environments.* We have enabled teams of robots to coordinate actions in close proximity. A particular accomplishment consists of an algorithm for coordinating a team of quadrotors equipped with cameras and IMUs providing the capability of vision-based cooperative control for aerial robots. (We emphasize this is based solely upon onboard sensing, without the aid of an external location system.)

9) *Indoor reconnaissance and search/rescue with time constraints.* Algorithms for high op-tempo indoor reconnaissance missions with n platforms covering m rooms in k minutes for a range of values of m, n, and k have been developed. This is accomplished through abstractions and algorithms for the rapid deployment of heterogeneous robot teams.

10) *VHF-dense urban networking*: We studied the communications properties when using the 40-MHz band in complex urban environments. We developed both theoretical and hardware-based experiments. We have worked on communication scenarios, protocols, and effective bandwidth measurements, and demonstrated same-channel, multi-user communications at very low power.

11) *Power-aware operation of teams*: While interested in high op-tempo deployment of teams of quadrotors in urban environments, we also advanced the energetics of deployment by benchmarking of component-level performance in quadrotor aircraft and performance across different length scales.

12) *Platform control in gusts*: We carried out experimental and simulation trade studies to quantify controller changes to improve estimation and adaption to gusts as well as trajectory generation for improved gust rejection. We addressed fast, online generation of motion plans for micro-UAVs that are informed by capabilities of UAVs subject to gusts.

13) *Landing and perching:* We addressed robust estimation, planning, and control for fast navigation in 3-D indoor environments, and used this to develop estimation and control techniques for fast perching and grasping.

4.6 Challenges for the Future

With decreases in the price/performance ratio of computation, communication, sensing, and storage technologies, we can imagine a new generation of inexpensive embedded, smart devices that are capable of autonomous operation—unmanned systems both in the air and on the ground. The opportunities afforded by this convergence of robotics and autonomy with the power of advanced computing, analytics, low-cost sensing, and new levels of connectivity will naturally lead to a networked team, where swarms of sensors, processors, and machines interact with and exchange information with humans. This new Internet-connecting of sensors, machines and humans addresses the design of devices that not only compute and sense, but also close the feedback loop and enable actuation in the physical world, thereby directly interfacing with the environment.

Making strides in this domain will require advances in autonomy (i.e., complex decisions will need to be made in an autonomous fashion due to the sheer scale and speed of the application domains). Big data will need to meet cyber-physical systems and robotics. We will need new architectures to distribute and manage local and cloud-based data and computation, new models for accessing the cloud for real-time operation while providing bounds on system performance, new algorithms to distribute computation and learning, and new abstractions for cloud services that are general enough to be used by a large number of applications.

We will also need to confront the inevitable challenges created by exponentially increasing complexity. Designing and deploying complex networked teams will require advances in five areas: heterogeneity, resiliency, learning, high-level task specification, and autonomous networking.

A more detailed description of our vision in these five areas is presented in the following subsections.

4.6.1 Heterogeneity

The consideration of a heterogeneous team adds a number of potential advantages, due to the robots' complementarity. In the case of ground and aerial vehicles, complementarity is given along two main axes: 1) mobility and 2) perception. For mobility, aerial vehicles have the ability to vary their vertical position, and hence, are capable of exploring complex 3-D environments, and do so efficiently.

Conversely, for the same sensor and processing payload, ground vehicles tend to be smaller in size than aerial vehicles and can potentially navigate through tighter spaces such as corridors and doorways. For perception, aerial vehicles can provide a birds-eye view of a scene, whereas ground vehicles can provide close-up ground perspectives.

This complementarity increases the complexity of the problem because some goals (or frontiers) can only be reached by one type of robot, while others can be reached by multiple types. If our aim is to seek paths that minimize the total time to execute a task in an environment, this gives rise to a hard combinatorial optimization problem. This problem is compounded by the fact that the environment is only partially known and as new parts of the environment are observed, the time optimality of previously generated plans no longer holds. This issue is unavoidable, unless the robots have prior knowledge about the environment or are able to predict what types of paths can be planned in the unobserved parts of the environment.

4.6.2 Resiliency

While there is much attention on designing software architectures when developing algorithms and writing software that makes individual nodes, the links between nodes, and even the network robust, we need to embrace "resilience" as a central paradigm for design and deployment. Resilience requires us to concurrently model and design the physical (spatial) network, the communication network, and the sensing network, and design algorithms that adapt to individual failures while rewiring connections across networks. As outlined previously, we lack a fundamental understanding of what it means to coordinate heterogeneous assets across multiple temporal and spatial scales. We also lack formal definitions of resilience that can be applied to dynamical systems that are physically constrained. Further, we need to understand how to translate global system-wide goals into resilient and robust local control policies. Indeed, our challenge is to create a basis of resilience and develop technical solutions that engender five key characteristics: 1) awareness, 2) diversity, 3) integration, 4) self-regulation, and 5) self-adaptation.

4.6.3 Learning

Our ability to execute tasks with an autonomous team is governed by a plethora of physical phenomena, such as ground interactions and shear deformations, that are either difficult to model from first principles or yield computationally intractable control systems. Though driving slowly enough may hide the difference between the models that underlie planning schemes and ground truth, this precludes any mission with an operational tempo or scale that necessitates high speeds. By abstracting away the complicated physical phenomena using statistical models of

the difference between expected and true state behavior, which we call the model disturbance, we may obtain the basic predictive power we need to drive a modern robotic system. The descriptive capability of the statistical model that captures unexpected events experienced by the platform is inherently contingent on the variation of the configuration space in which the team operates.

In practice, we might not have access to a data set for a particular operational environment or the environment may have such dramatically different surfaces that a generic model is too imprecise. In such cases, an adaptive approach, where one sequentially revises estimates of the model disturbance based on new information, is advantageous. The solution that we envision for this online refinement process is the development of statistical learning models to describe the space in which the team is operating. We envision the use of sensory input to feed into these learning schemes that aim to adapt control actions to the observed environmental uncertainty.

4.6.4 High-Level Task Specification

MAST has advanced the state of the art to the point where the execution of simple primitives can be considered resolved. It is possible for us to specify a configuration that the team must adopt or ask the team to locate and land on a target. In future scenarios, we think that it will become necessary to specify tasks at a higher level. The two intellectual challenges that we envision are 1) transformation of high-level specifications into low-level constraints that are specified to the team and 2) modification of existing algorithms so that they are targeted toward satisfiability problems instead of achievability problems. Our vision is that team operators assign the team a task specification such as "secure a 50-yd perimeter around the operation center". Algorithms then translate this specification into the identification of points that have to be secured and decide on the deployment of land and air assets to accomplish the task. Software architectures that allow us to automatically map such tasks onto the physical and information assets and compile them into behaviors must also be considered as we think about questions of specification and design.

4.6.5 Autonomous Networking

The previous four key characteristics inherently rely on communications among autonomous nodes, sensors, knowledge bases, and humans. MAST has definitively shown that viewing networking as a separate component leads to distinctly suboptimal results, and that networking is a critical aspect of all the previous desired characteristics. Autonomous networking should exploit cognitive radio techniques and take advantage of the heterogeneity across power and complexity. Resiliency is needed not only with respect to collaborative ensemble goals, but also

with respect to the communications needed to achieve those goals. Adaptation and collaboration will be key to mission effectiveness, especially in the face of adversarial deployment of similar systems.

4.7 CNC Task Summary

CENTER LEAD: Vijay Kumar - University of Pennsylvania			
ARL TECH LEAD: Brian Sadler			
Project #	**Project Title**	**PI**	**ORG**
CNC-16-1	Novel Radios Sub-Thrust	Sarabandi	Michigan
CNC-16-1.1	Short Range HF Communication in Highly Cluttered Environments	Sarabandi	Michigan
CNC-16-2	Communication and Control Software Sub-Thrust	Ribeiro	Pennsylvania
CNC-16-2.1	Distributed information games for autonomous team behavior	Ribeiro	Pennsylvania
CNC-16-2.2	Communication Coverage for Situational Awareness	Rus	MIT
CNC-16-2.3	Exploiting Heterogeneity in Wireless Aerial and Ground Robotic Networks	Fierro	Mexico
CNC-16-2.4	Robust estimation, planning and control for fast navigation in 3-D indoor environments	Kumar / Daniilidis	Pennsylvania
CNC-16-3	Architecture and Command/Control Sub-Thrust	Kumar	CMU
CNC-16-3.1	Dynamic Goal Assignment and Trajectory Generation for Cooperative Teams of Micro Air Vehicles Exploring Complex 3D Environments	Michael	CMU
CNC-16-3.2	Vision Based Cooperative Control for Aerial Robots	Kumar / Daniilidis	Pennsylvania
CNC-16-3.3	Rapid deployment of heterogeneous robot teams: abstractions, algorithms and experimentation.	Carpin / Tanner / Kumar	UC Merced / UDel / UPenn

4.8 Section Bibliography

Ayanian, N., Kumar, V., and Koditschek, D. (2011). Synthesis of controllers to create, maintain, and reconfigure robot formations with communication constraints. In Robotics Research (pp. 625-642). Springer Berlin Heidelberg.

Bhattacharya, S., Michael, N., and Kumar, V. (2013). Distributed coverage and exploration in unknown non-convex environments. In *Distributed autonomous robotic systems* (pp. 61-75). Springer Berlin Heidelberg.

Choi, J., Dagefu, F. T., Sadler, B. M., and Sarabandi, K. (2016). Electrically small folded dipole antenna for HF and low-VHF bands. *IEEE Antennas and Wireless Propagation Letters*, *15*, 718-721.

Dagefu, F. T., Choi, J., Sheikhsofla, M., Sadler, B. M., and Sarabandi, K. (2015). Performance assessment of lower VHF band for short-range communication and geolocation applications. *Radio Science*, *50*(5), 443-452.

Dagefu, F. T., Verma, G., Rao, C. R., Paul, L. Y., Fink, J. R., Sadler, B. M., and Sarabandi, K. (2015). Short-range Low-VHF channel characterization in cluttered environments. *IEEE Transactions on Antennas and Propagation*, *63*(6), 2719-2727.

Fink, J., Ribeiro, A., and Kumar, V. (2012). Robust control for mobility and wireless communication in cyber–physical systems with application to robot teams. *Proceedings of the IEEE*, *100*(1), 164-178.

Fink, J., Twigg, J., Yu, P., and Sadler, B. M. (2013). A parsimonious model for wireless connectivity in robotic networks. In *Global Conference on Signal and Information Processing (GlobalSIP), 2013 IEEE* (pp. 855-858). IEEE.

Lin, D. T., Chae, H., Li, L., and Flynn, M. P. (2013). A low-power adaptive receiver utilizing discrete-time spectrum-sensing. *IEEE Transactions on Microwave Theory and Techniques*, *61*(3), 1338-1346.

Lin, D. T., Li, L., Farahani, S., and Flynn, M. P. (2012). A flexible 500 MHz to 3.6 GHz wireless receiver with configurable DT FIR and IIR filter embedded in a 7b 21 MS/s SAR ADC. *IEEE Transactions on Circuits and Systems I: Regular Papers*, *59*(12), 2846-2857.

Stephan, J., Fink, J., Charrow, B., Ribeiro, A., and Kumar, V. (2014). Robust routing and multi-confirmation transmission protocol for connectivity management of mobile robotic teams. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 3753-3760).

Twigg, J. N., Fink, J., Paul, L. Y., and Sadler, B. M. (2013). Efficient base station connectivity area discovery. *The International Journal of Robotics Research*, 0278364913488634.

Short Range HF Communication in Highly Cluttered Environments

Task: CNC-16-17-1.1

Kamal Sarabandi, Jihun Choi
Dept. of Electrical and Computer Engineering
The University of Michigan
Ann Arbor, USA
saraband@umich.edu

Brian M. Sadler, Fikadu Dagefu
Army Research Laboratory
Adelphi, USA
brian.m.sadler6.civ@mail.mil

***Abstract*—We demonstrate bandwidth enhancement of a highly miniaturized low-VHF antenna via a non-Foster matching technique for low-power, short-range, wireless transmission systems. A recently developed miniature antenna ($0.013\,\lambda_0 \times 0.013\lambda_0 \times 0.02\lambda_0$ at 40 MHz) for this application is shown to have much higher radiation efficiency compared to similar-sized antennas, but its bandwidth is rather narrow. In order to improve the bandwidth while maintaining the best possible efficiency for the given size of the passive antenna, we introduces a non-Foster matching network realized by a negative impedance converter (NIC). We also presents a trade-off study between degree of bandwidth enhancement and the stability of the NIC circuit. Having a power efficiency advantage, a fabricated non-Foster matched antenna achieves a fractional bandwidth of about 2.15 %, an improvement of more than twice the bandwidth of the passive version. We also propose a minimalistic design for node localization based on a small number of access points. Our method uses the received signal strength (RSS) gradient as a direction of arrival (DoA) estimator. Though non-uniform shadowing and multipath adversely affect RSS, these are much less at lower VHF. We show through high fidelity simulations that our approach has promise as a rapidly-deployable localization system. Furthermore, we investigate a time of arrival based localization technique for nodes deployed in complex GPS-challenged environments. The proposed technique exploits spatial and polarization diversity via anchors equipped with polarization diversity antennas operating at low frequencies. Empirically derived effective dielectric constants are used for penetration delay compensation. Our full-wave simulations of a complex scene suggest that sub-meter localization error can be achieved, with low-power transmission at upper HF and lower VHF bands.**

***Index Terms*—Bandwidth Enhancement, HF band, low-VHF Band, localization, non-foster matching technique, polarization diversity, received signal strength, small antennas**

I. Introduction

Reliable communication is of paramount importance in many military applications. In scenarios where the communication environment is complex, such as urban or indoor environment, signal attenuation and multipath poses stringent conditions on power and signal processing. This situation is encountered by MAST small robotic platforms in all three scenarios under consideration including: 1) Urban search and mapping, 2) Cave or demolished building search and mapping, and 3) Perimeter defense. Radio communication is far more complex for such small robotic platforms since important wireless infrastructure, such as base stations with predetermined proper positions, are absent and high transmit power is not available. Much effort has been devoted over many decades to design modulation schemes appropriate for different communication bands to overcome challenges introduced by the communication channel and factors such as noise and interference. In general it is understood that channel complexity is far less at lower frequencies since significant signal penetration through/over obstacles is possible and reflection, scattering, and diffraction is far lower than what is encountered at upper UHF and higher bands. High penetration through soil and building materials allows for direct communication through caves and multiple layers of walls and buildings. Reduced reflection, scattering, and diffraction will be responsible for eliminating severe signal fluctuations, and fading which usually put stringent requirements on power, system bandwidth, as well as waveform and modulation. The severe signal fluctuations and intermittent signal drop-off that are unavoidable at higher frequencies are significantly reduced at these lower frequencies due their reduced reflection and scattering. It should be mentioned, however, that multipath effect is the main source of connectivity between a transmitter and a receiver in a complex environment at higher frequencies, where line-of-sight signal is extremely weak. On the other hand, at lower frequencies, the size of the conventional antennas is very large, making application of such systems limited to stationary nodes or, in the case of mobile nodes, the transmit power has to be high enough to compensate for the poor efficiency of the antennas. With advances in the design of small antennas, a careful analysis of propagation and scattering at the HF range using accurate physics based simulations as well as measurements in realistic scenarios is vital. Based on this analysis, the potential of using the upper HF band for applications in HF *ad hoc* networks can be assessed. Such a system can provide long distance coverage unhampered by walls, buildings, and foliage albeit being limited by instantaneous bandwidth. As mentioned before the limited bandwidth may not be a serious issue in the absence of fading.

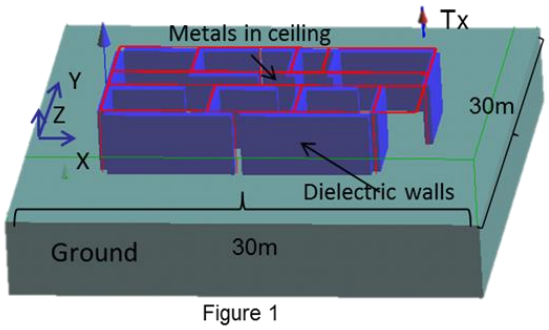


Fig. 1. Example of a single story office building with metallic frame.

With proper design of miniaturized antennas, it is possible to achieve multi-band antennas without significantly increasing their size.

II. RESEARCH ISSUE AND TECHNICAL APPROACH

A. Research Issue

Physics-based indoor propagation analysis in the HF band complemented by measurement results in complex indoor scenarios as well as the development of extremely small, light-weight HF antennas having reasonable gain and multi-band capability will play a vital role in supporting reliable direct communication among MAST platforms in complex environments including caves, underground, and through multiple layers of walls or buildings. Combining such communication nodes to perform beam-forming in an *ad hoc* manner among distributed nodes of MAST platforms allows very low power operation in very difficult propagation environments. The latter will be accomplished in collaboration with algorithms developed by Prof. Madhow of UCSB and Prof. Raghu Mudumbai of U. Iowa.

B. Technical Approach

As mentioned earlier, to study the pros and cons of an *ad hoc* wireless system at upper HF band, we will need to study the wave propagation behavior in complex propagation environments such as indoor and urban settings. In such environments building walls, ceiling and floors are the dominant scatterers. Fig. 1 shows a model of a complex building structure made from concrete cinder blocks where a transmitter is positioned outside the building (ceiling not shown). The ground is modeled as a half-space dielectric medium while the walls are modeled as thin dielectric slabs. A concrete ceiling is also included in the model. Realistic dielectric values are used for the walls, ground and ceiling. To make the environment more complex, a metallic frame is considered as support for the dielectric walls and ceiling. A short dipole located outside is utilized as a Tx antenna and field coverage inside and outside the building is computed as a function of frequency by utilizing accurate solutions. For such building walls, electromagnetic waves at high frequencies (UHF band and higher) will experience significant scattering. Basically, each ray that hits the walls will scatter in multiple directions. Each of the reflected

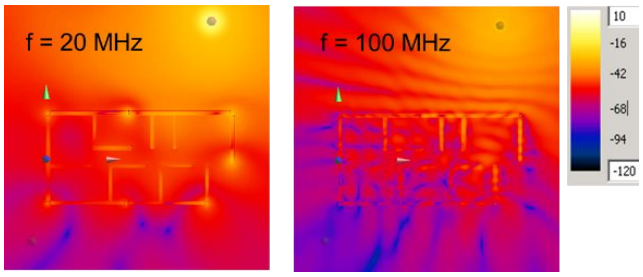


Fig. 2. Field coverage on a horizontal plane which is 1.5 m (=λ/10) above the ground. The left and right figures show, respectively, the field coverage at 20 MHz and 100 MHz, making communications at such low frequencies more reliable.

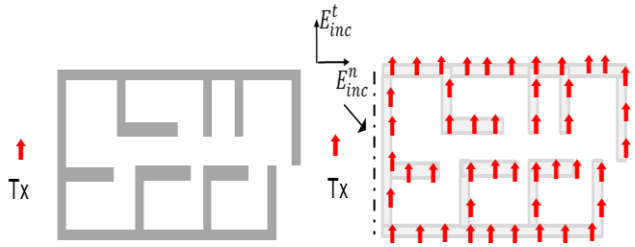
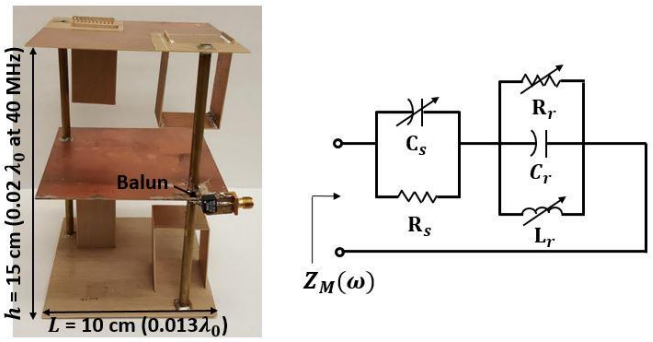


Fig. 3. The Proposed HF propagation model utilizes Rayleigh-Gans approximations in conjunction with a near-ground propagation model to take into account scattering from building walls. Polarization currents are used to capture the effect of the walls.

rays will bounce off of a different wall and multiply into many other rays, creating a highly fluctuating field distribution within the building. At lower frequencies (VHF and lower) models based on ray tracing do not work. For small buildings full-wave analysis methods such as the finite difference time domain (FDTD) method can be implemented on high performance computers. The signal coverage on a plane parallel to the ground at a height z_o = 1.5m is plotted for the building (see Fig. 1) based on a FDTD full-wave simulation for two different frequencies (20MHz and 100MHz) as shown in Fig. 2.

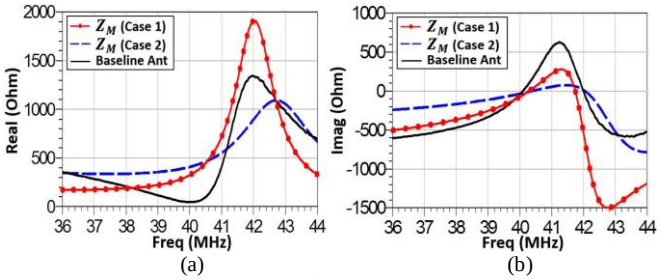
As expected, when a vertical dipole is used as Tx, the effect of the vertical metallic cylinders (found in walls in some non-residential buildings) can be seen in the field coverage at both frequencies. However, the horizontal metallic cylinders are not significant since the radiated field is vertically polarized. It is clear that even at 100MHz there is significant multipath and attenuation. At 20MHz, however, the walls appear as thin, effectively homogeneous, dielectric slabs with small reflectivity and high transmissivity. Under such conditions multi-path is minimized and the signal can propagate through walls with little attenuation. Another advantage is that the free-space path-loss, which increases with the square of frequency, is lower at lower frequencies. We also need to consider issues for near-ground propagation.

For transmitters and receivers near the ground the field decays as $1/R^2$ and the coefficient is proportional to the relative height (compared to the wavelength) of the transmitter and receiver. However, this degradation is not as important as the other two factors. We will examine different propagation scenarios for estimation of path loss using full-wave solutions for medium-size indoor scenarios and develop an efficient HF indoor and outdoor propagation model to analyze relatively large indoor and outdoor settings.



(a) (b)

Fig. 4. Baseline miniature passive antenna and its impedance model $Z_M(\omega)$.



(a) (b)

Fig. 5. Impedance model examples for the baseline antenna: (a) real and (b) imaginary parts of $Z_M(\omega)$ together with measured response of the baseline antenna.

The incident fields on the walls that have line-of-sight (LoS) with the transmit (Tx) antenna will be calculated by taking into account the near-ground effects. The fields will then be converted to equivalent currents using the Rayleigh or Rayleigh-Gans formulations, together with the current on the Tx antenna, will be used to illuminate the other non-line-of-sight (NLoS) walls, and calculate the field coverage inside and outside the building structure (see Fig. 3). Once the fields in the building are calculated, we will use equivalent currents and the analytical wave approach to calculate the field outside and over long distances. Parameters of interest include attenuation and field coverage, field fluctuation statistics, and coherence bandwidth for multiple MAST scenarios.

We will also investigate issues related to design of miniaturized antennas with enhanced bandwidth for operation at upper HF frequencies. Recently, we have designed and demonstrated performance of an extremely small antenna with dimensions as small as λ/100 by λ /100 by λ /300. The proposed antenna was not only small in size but also has very small weight (~15 g) since it was designed without any dielectric substrate unlike many of the state-of-the-art antennas. This antenna is capable of radiating purely vertical polarization with a relatively high gain of -18dBi. Investigations will be carried out to make improvements on this design, by increasing its bandwidth and efficiency. Tradeoff studies related to the center frequency, antenna size, data rate, and power consumption will also be carried out.

To tackle the challenge associated with the relatively small bandwidth of extremely small antennas such as the one described above, the use of multi-band antennas will be pursued. One possible approach is to design an array of a few of the miniaturized antennas operating at HF frequencies close to each other (e.g. 20 MHz, 24MHz, 28MHz) with one of the design goals being minimizing the overall size of the array. By utilizing the center frequencies as well as the first few harmonics, the available frequency band could be increased dramatically. The use of closely positioned narrowband antennas to increase bandwidth by taking advantage of mutual coupling between antenna elements is another approach that will be investigated.

III. RESEARCH ACCOMPLISHMENTS

A. Bandwidth Enhancement of the Miniature Low-VHF Antenna

In [1]-[3] the non-Foster matching technique is applied to unmatched small monopole or loop antennas. This technique utilizes a single negative capacitor and/or a single negative inductor produced by a negative impedance converter (NIC) to cancel out the reactive components of the antennas over a wide frequency range. However, for our baseline antenna (Fig. 4(a)), reactance cancellation is more challenging due to the steep impedance variation near the resonant frequency. Our goal is to broaden the effective bandwidth of the baseline passive antenna design while providing the best possible radiation efficiency for a given form factor. An advantage of this approach is that a single NIC can be applied directly to the antenna Balun, enhancing the overall power efficiency compared to using two NICs as in [1]. For its successful realization, challenging issues in antenna impedance modeling and stability with respect to system sensitivity are involved, which control the degree of bandwidth enhancement.

The extent of bandwidth enhancement by means of the proposed non-Foster matching technique critically depends on the accurate representation of the antenna impedance model. This will allow design of a NIC circuit to ameliorating the impedance variation of the antenna around its resonance. The baseline passive antenna has relatively narrow bandwidth (the 3 dB bandwidth of this design is 0.35 MHz or 0.87 %), which indicates a strong dependence of input impedance on frequency. Usually antenna impedance near the resonance is simply modeled by an RLC circuit. To improve the bandwidth by a significant amount a simple RLC circuit model is not effective. A more appropriate model of the antenna at hand is shown in Fig. 4(b), where impedance model $Z_M(\omega)$ consists of two series parallel circuits and can be expressed as

$$Z_M(\omega) = R_r \parallel j\omega L_r \parallel \frac{1}{j\omega C_r} + R_s \parallel \frac{1}{j\omega C_s}, \qquad (1)$$

where ω is the angular frequency. Fig. 5 illustrates examples of the real and imaginary parts of $Z_M(\omega)$, together with the measured antenna impedance. Case 1 parameters are $C_S = 6\,pF, R_s = 3.5\,k\Omega,\ R_r = 1.8\,k\Omega, L_r = 239\,nH,$ and $C_r =$

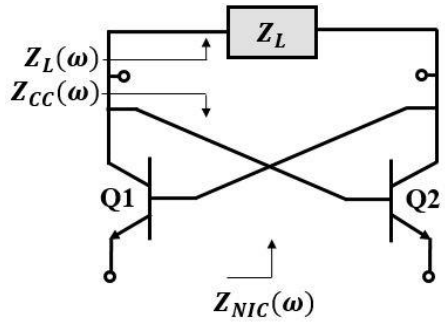


Fig. 6. A schematic of a cross-coupled pair consisting of two NPN bipolar junction transistors (BJT). $Z_L(\omega)$ is the impedance to be inverted with its negative value.

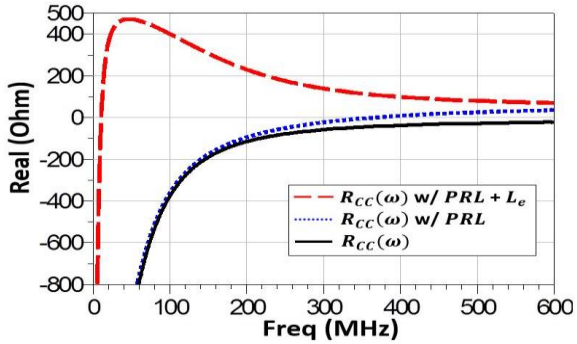


Fig. 7. The real part of the cross-coupled pair in Fig. 6 is shown, which is dependent on a parallel inductor-resistor *PRL* along the cross-coupled path and an inductor L_e in series with a resistor at emitters of the BJTs.

$60\ pF$. Case 2 parameters are $C_S = 10\ pF, R_s = 810\ \Omega, R_r = 880\ \Omega, L_r = 204\ nH$, and $C_r = 68\ pF$.

Note that the modeling parameters affect both the degree of bandwidth enhancement and system stability. For example, comparing the imaginary part of $Z_M(\omega)$ for the two cases shows that Case 1 better represents the baseline antenna. Therefore, an NIC circuit can be designed for better reactance cancellation through the NIC, but its faster resistance variation with frequency may lead to a stability issue with the non-Foster matching network (this will be discussed in the following section). This illustrates that the modeling parameters must be optimally chosen to achieve the maximal bandwidth under stable performance of the system.

A Bipolar Junction Transistor (BJT) cross-coupled pair (see Fig. 6) is widely used to produce a negative resistance in oscillator design or to function as a NIC. However, a careful stability analysis of the NIC is needed for non-Foster matching applications, due to the internal positive feedback that can lead to circuit oscillation. We design the matching circuit with the NIC to meet the stability criterion described below, simulate the circuit, and apply the Nyquist stability test to verify the validity of the design. The first step in this process is to characterize the impedance of the cross-coupled pair $Z_{CC}(\omega) = R_{CC}(\omega) + jX_{CC}(\omega)$ with a DC biasing circuit as a function of frequency in the small-signal regime. Since the cross-coupled pair itself generates a negative resistance, $R_{CC}(\omega) < 0$ when a passive load impedance $Z_L(\omega) = R_L(\omega) + jX_L(\omega) = \infty$ is an open

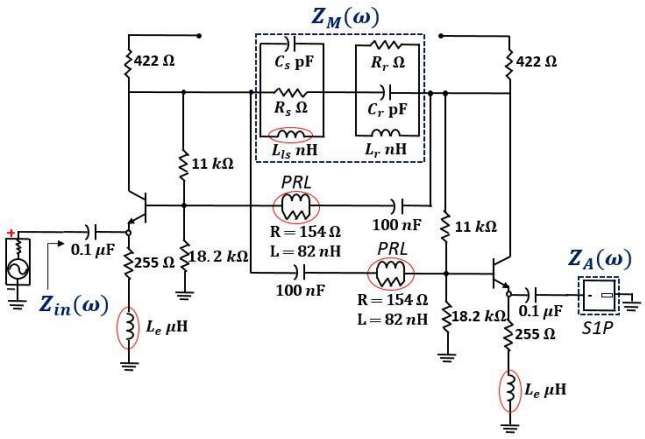


Fig. 8. A schematic of the non-Foster matching network designed to ensure stable system performance.

circuit (this is plotted as a black dashed line in Fig. 7). Also, to avoid oscillation conditions when $Z_L(\omega)$ is connected, the net resistance of the network must be positive, expressed as

$$Re[Z_{CC}(\omega) + Z_L(\omega)] > 0 \ or \ |R_{CC}(\omega)| < R_L(\omega). \qquad (2)$$

Here, $Z_L(\omega)$ corresponds to the aforementioned $Z_M(\omega)$ for non-Foster matching of the baseline antenna. In order to satisfy (2) at all frequencies, we proceed as follows.

To make $R_{CC}(\omega)$ positive at higher frequencies, the feedback factor along each cross-coupled path is adjusted by adding a parallel inductor-resistor *PRL* ($124\ \Omega \parallel 82\ nH$) in series with a feedback capacitor. Furthermore, an inductor L_e ($10\ \mu H$) in series with the emitter resistor is added for each BJT. With the added components, $R_{CC}(\omega)$ becomes positive above 12 MHz as depicted in Fig. 7. To fulfill (2) at lower frequencies, $Z_M(\omega)$ is modified by adding an inductor L_{ls} in parallel with R_s and C_s to produce high impedance below 12 MHz.

Fig. 8 illustrates a schematic of the designed matching network and the relevant components added to the design in Fig. 6 are circled. Here, $Z_{in}(\omega)$ represents the sum of the impedance of the floating NIC, $Z_{NIC}(\omega)$ (see Fig. 6), and the measured antenna impedance, $Z_A(\omega) = R_A(\omega) + jX_A(\omega)$. The total loop resistance must also be positive in order to avoid circuit oscillation. This design was simulated over a range of values for the elements in $Z_M(\omega)$, taking into account the component tolerances and reduction of the antenna impedance variation over the frequency range of interest while avoiding oscillations. Next, we conduct a large-signal S-parameter simulation to ensure overall robustness to nonlinear effects of the BJT with respect to large AC signals and variations in power supply values.

Based on the design described above, a non-Foster matching network was fabricated and characterized. In order to fine tune $Z_M(\omega)$ in the presence of antenna impedance variations arising from the measurement environment (e.g., proximity effects) and component and fabrication tolerances, mechanically adjustable elements are utilized (they are marked with arrows

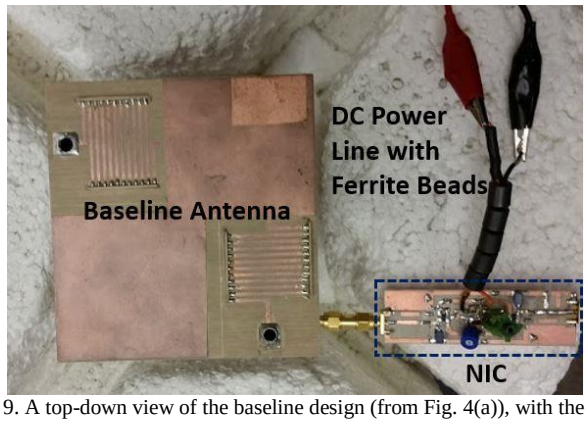


Fig. 9. A top-down view of the baseline design (from Fig. 4(a)), with the NIC connected directly to the antenna Balun.

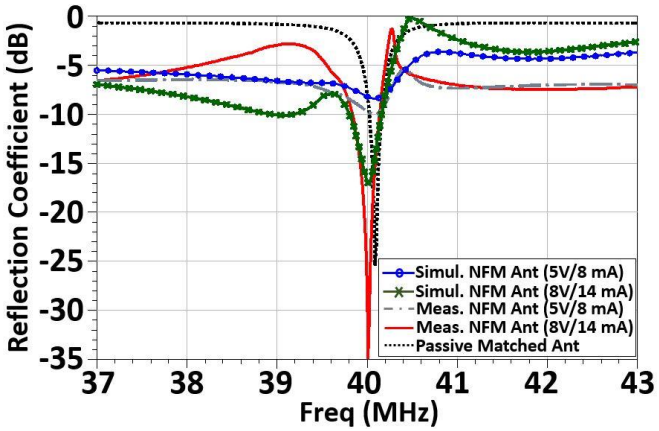


Fig. 10. Simulated and measured reflection coefficients as a function of power supply value expressed as voltage to current ratio.

in Fig. 4(b)). The fabricated NIC is connected directly to the antenna Balun along the null plane bisecting the antenna structure (see Fig. 4 and 9). This ensures that introducing the NIC does not affect the characteristics of the passive antenna such as its input impedance and radiation pattern [4]. Multiple ferrite beads are used on the DC power line for high frequency noise and parasitic suppression. They are also utilized on the coaxial feed cable to suppress common mode current flowing on the outer conductor of the coax.

The fabricated non-Foster matched antenna is characterized using a vector network analyzer. Large-signal measurements as a function of input power and power supply voltage are performed to examine nonlinear behaviors of the BJTs in the NIC. We first fix the input power to 7 dBm, corresponding to a relatively low power transmission in the low-VHF band. Keeping the input power fixed, we tested with different voltage/current ratios, measured reflection coefficients of the non-Foster matched antenna as a function of frequency, and compared them with circuit simulation results as shown in Fig. 10.

We also measured the transmission coefficient for the same cases as in Fig. 10. For this measurement, a dipole with a length of $0.2\,\lambda_0$ at 40 MHz is used as a receiving antenna. The transmission coefficients as a function of frequency are

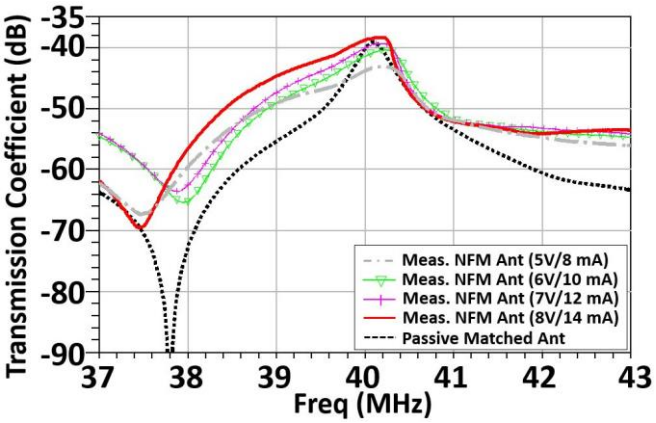


Fig. 11. Measured transmission coefficients as a function of power supply value. A $\lambda_0/5$ dipole is used as a receiving antenna.

TABLE I
COMPARISON OF ANTENNA MATCHING PERFORMANCE

Matching Type	3 dB BW	η_f (within the BW)
Passive	0.87 %	·
Non-Foster	2.15 %	1.3 to 5.9

depicted in Fig. 11.

The measurement results show the dependence of the impedance matching and DC bias. Note that, when the input AC signal is small compared to the DC bias, the cross-coupled pair properly functions as a NIC in the small-signal regime where balanced operation of the BJTs in the pair is maintained. This is why the matching performance is better for higher DC bias values as can be seen in the results. However, over the DC bias range indicated in Fig. 10, the cross-coupled pair turns into a regenerative circuit that gives rise to matching degradation and system instability. A similar behavior was also observed by sweeping the input AC signal while keeping a fixed DC bias (not shown here). Thus, the maximal bandwidth of the matching network under stable operating conditions can be determined with a proper DC bias for a given input power.

Since the non-Foster matched antenna requires DC power to operate the NIC, the overall power efficiency should be considered to evaluate the effectiveness of the proposed approach against the baseline passive design. We define a power efficiency factor η_f as

$$\eta_f = \frac{|S_{21(AM)}|^2}{|S_{21(PM)}|^2} \cdot \frac{P_{DC(PM)}}{P_{DC(AM)}}. \quad (3)$$

Here, $S_{21(AM)}$ and $S_{21(PM)}$ are measured transmission coefficients with active and passive matching, respectively. $P_{DC(PM)}$ is the DC power for transmission in the passive case, and for the active case $P_{DC(AM)}$ is the DC power for transmission plus the power for the active matching circuit. Aligning with low-VHF radio experiments, P_{DC} for the transmitter is set to be 3.3 V/ 240 mA, and P_{DC} for the NIC is

chosen to be 8 V/ 14 mA. With these values, at a 40 MHz center frequency, the active matching enhances the 3 dB bandwidth from 0.87 % to 2.15 % with a power efficiency advantage such that $\eta_f > 1$ over the entire bandwidth. The performance comparison between the passive and active matching is summarized in Table I.

B. Localization via the Received Signal Strength Gradient

Localization and tracking of targets are critical objectives in a variety of fields, e.g., in radar and mobile robotic systems. Use of global positioning system (GPS) is not always feasible due to cost and/or poor penetration. Many localization techniques rely on time of arrival (ToA), angle of arrival (AoA), or received signal strength (RSS). ToA-based methods based on time-of-flight of a wireless signal suffer from errors due to clock drift and variable processing delays. Classical AoA-based approaches rely on array-processing, requiring multiple synchronized antenna elements per node, and have many challenges, e.g., the location estimate is very sensitive to multipath. Classical RSS-based localization approaches attempt to estimate distance-to-source from the received power; this requires estimating various nuisance parameters (e.g., path loss exponent) and in complex scenes, RSS is corrupted by multipath and shadowing [5]. In this work, we take a fundamentally different approach; we use RSS to compute the direction of arrival (DoA) to fixed-location nodes (beacons), and fuse the DoA's to localize the node. This method relies on the basic fact that under ideal propagation, the RSS spatial gradient points towards the source. At microwave frequencies, huge fluctuations in RSS occur over small spatial scales due to small-scale fading. However, these effects are much less prevalent at lower VHF. In addition, we can employ statistical methods to greatly reduce the effect of RSS outliers. Our framework yields point estimates of the DoA as well as uncertainty regions. RSS is easy to measure, does not require an array to process, and since the gradient is a relative measure (a slope), it requires no calibration. Estimation of the RSS spatial gradient does require multiple spatially separated RSS measurements.

Our model for DoA estimation from RSS data is described in [6]. The RSS surface is locally modelled as a linear function of sampling location, and the gradient of the model is a DoA estimate. Notationally,

$$y = X\beta_i + \varepsilon \tag{4}$$

where y_i is a $N \times 1$ vector of RSS values of beacon i at N sampling locations, X is a N × 2 matrix of those RSS sampling locations, and $\varepsilon = N(0, \sigma^2 I_N)$, where I_N is the N × N identity matrix. We assume the presence of $B \geq 2$ beacons with known fixed locations. Thus, a node seeking to localize itself measures the RSS with respect to each of the B beacons over X. We adopt a Bayesian model over 1 which enables us to incorporate prior information on σ^2 and to reason probabilistically about the

Algorithm 1 RSS gradient-based localization

for m=1 to M **do**
  for i=1 to B **do**
    $\boldsymbol{\beta_i} \leftarrow \boldsymbol{p}(\boldsymbol{\beta_i}|\boldsymbol{y}, \boldsymbol{X})$
    $\boldsymbol{n_i} \leftarrow \begin{pmatrix} \mathbf{0} & \mathbf{-1} \\ \mathbf{1} & \mathbf{0} \end{pmatrix} \frac{\boldsymbol{\beta_i}}{\|\boldsymbol{\beta_i}\|}$
    $\boldsymbol{\omega_i} \leftarrow \|\boldsymbol{\beta_i}\|$
  end for
  $\boldsymbol{l_m} \leftarrow \left(\sum_i \boldsymbol{\omega_i} \boldsymbol{n_i} \boldsymbol{n_i^T}\right)^{-1} ((\sum_i \boldsymbol{\omega_i} \boldsymbol{n_i} \boldsymbol{n_i^T} \boldsymbol{b_i})$
end for

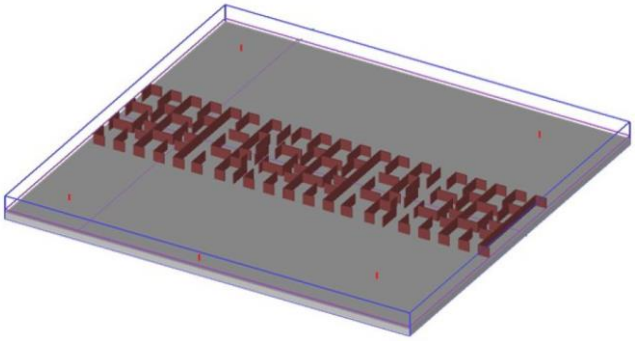


Fig. 12. FDTD simulation of an office building with realistic materials. All nodes are at assumed height of 0.1λ and at center frequency of 40 MHz. The total scene is 122 m × 115 m. Beacons (red dots) with known locations surround the building.

DoA estimate β_i, where i ranges from 1 to B. We can use previous studies characterizing the shadowing variance at lower VHF to derive distributions over σ^2. We use $p(y|X, \beta) = N(X\beta, \sigma^2 I_N)$, a Normal distribution. Using standard statistical theory, we can derive the posterior probability distribution $p(\beta_i|y, X)$ over each β_i.

We next combine $\{\beta_i\}_{i=1}^{i=B}$ into a localization estimate. In the general case, the node's orientation θ is unknown (with respect to the global coordinate system). θ is straightforwardly estimated; but for simplicity in the sequel we assume an estimate of it is known (e.g., via compass), so that the β_i are represented in the global coordinate system. Our localization algorithm is presented in Algorithm 1. For each of M Monte Carlo iterations, we sample all B direction vectors from their posterior distributions. We apply a weighted least squares procedure on these DoA's to obtain a point estimate of the node location l based on known beacon locations bi. M such Monte Carlo runs yield an empirical distribution over node location. We must account for correlations and outliers in the data to ensure the estimates are meaningful. Positive spatial correlation arises when nearby RSS samples encounter correlated shadowing. Without correction, the true uncertainty in the DoA (and hence localization) is underestimated. Outliers arise from fading and non-uniform shadowing, and bias the estimated DoA from the true source direction. We use standard statistical procedures for correcting spatial correlation and outliers; details will be described elsewhere.

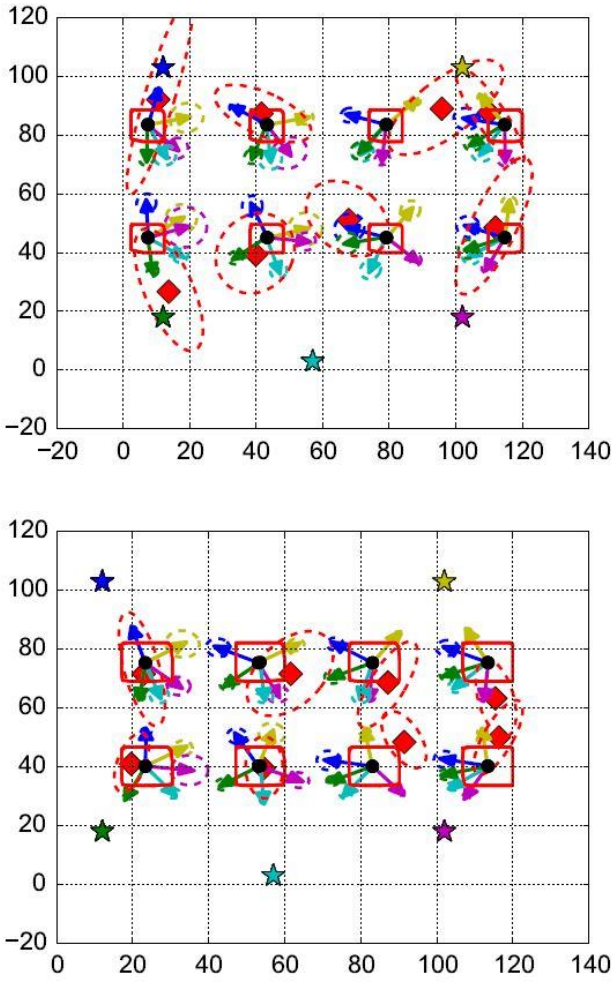


Fig. 13. Localization results for a mobile agent which samples 8 different sampling regions within the office scene in Fig. 12 of sizes (top) ≈ 10m × 10m and (bottom) ≈ 12m × 12m. Figure elements: beacon (colored star); estimated DoA based on RSS (arrow colored to match corresponding beacon); DoA uncertainty region (colored dashed ellipse encircling DoA arrow); convex hull of the sampling region (rectangular red region); true location (i.e., centroid of sampling region) (black dot); location estimate (red diamond); location uncertainty ellipse (dashed red ellipse). Top: Median localization error is 10:1 m. Bottom: Median localization error is 7:4 m.

We have conducted high-fidelity simulations at lower VHF using a finite-difference time-domain (FDTD) full wave solver in a complex indoor office environment shown in Fig. 12. Fig. 13 shows the results; the caption has further details. Note that as the RSS sampling region grows, the typical localization error decreases and the uncertainty ellipses shrink.

C. Full-Wave Analysis of Time of Arrival Based Localization with Polarization Diversity

Accurate localization of mobile nodes in GPS-challenged scenarios has various civilian and military applications including navigation of robotic platforms deployed in complex infrastructure-poor environments for situational awareness as well as search and rescue missions. Conventional techniques that operate at microwave frequency bands either have poor accuracies or require extensive training and/or site specific measurement database and often assume the availability of infrastructure. For near-ground nodes deployed in such environments, recent studies suggest that lower frequencies (20-120 MHz) offer significant advantages due to the presence of a dominant line-of-sight (LoS) like path even when there are a multitude of scatterers between the nodes. This is due to the improved penetration and reduced multipath levels [7]. The presence of a dominant direct path makes the near-ground low frequency channel a good candidate to improve the accuracy of time of arrival (ToA) and array based localization techniques.

However, the presence of electrically large conductive scatterers creates signal components that will introduce significant errors. The strength of the scattered components is a function of the electrical size, material characteristics, frequency, range and polarization. For example, a vertically oriented metallic cylinder will reflect a vertically polarized wave, but minimally affects horizontally polarized field, and vice versa. This suggests that, depending on the orientation of nearby scatterers, one polarization may be preferred to minimize scattering and preserve a dominant LoS-like channel. Another important factor is phase aberration due to penetration delays through obstacles, which can significantly bias the ToA based range estimate, resulting in a biased localization estimate. Here, we report our initial investigation of a ToA based localization technique that exploits polarization diversity and empirical penetration delay compensation.

We consider a ToA based localization technique that leverages the low frequency near-ground channel in conjunction with polarization diversity and penetration delay compensation to significantly reduce ranging and localization errors. A typical scenario of interest consists of a minimum of three anchors with known locations as well as a mobile node deployed in an indoor or urban scenario. We assume all nodes are equipped with co-located polarization diversity antennas and the anchors transmit known Gaussian pulses in all three antenna orientations. The mobile node captures the three components of the signal from the various anchors. We assume orthogonal transmission among the anchors, e.g. using code division multiple access. As alluded to in the previous section, although the near-ground low frequency channel has a LoS-like dominant component due to better penetration and reduced multipath [7], the presence of electrically large scatterers near the antennas can cause some polarizations to experience significant multipath depending on the orientation of the dominant obstacle.

For a given location of the mobile node, the received signal from each anchor will have three components corresponding to the three orientations of the antennas. The receiver employs a matched filter for each polarization, matched to the Gaussian pulse. Three different ToA estimates are then computed by taking the difference between receive pulse time for each polarization and reference transmit time. We then choose the smallest ToA as our estimate. This ensures that the polarization that experienced the least multipath is selected as the ToA estimate for that anchor.

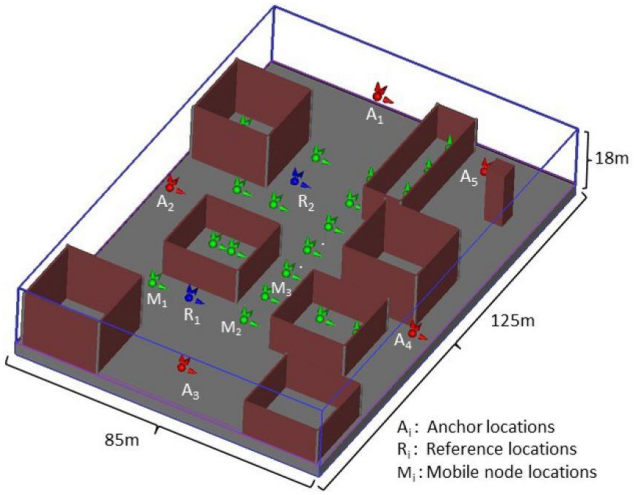


Fig. 14. An indoor-outdoor scenario consisting of dielectric and metallic scatterers with 5 anchors, 2 reference locations and 29 mobile node locations.

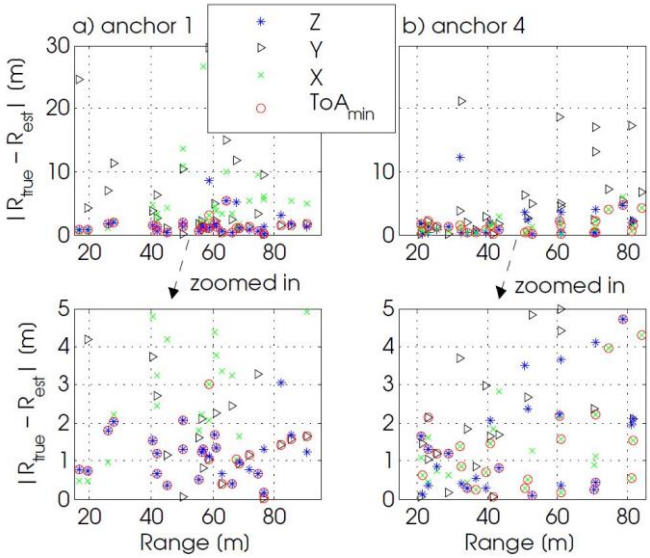


Fig. 15. The ranging errors based on ToA estimated from each polarization pairs are shown for anchors 1 and 4. The final estimate (denoted byToA_{min}) is computed by taking the smallest ToA estimate based on all 3 polarizations; note that the final choice varies with node location.

In the presence of multiple layers of large scatterers, even if there is a dominant direct path, the spatially varying penetration delay through multiple layers of scatterers could be significant. To combat this, an empirically estimated anchor-dependent effective dielectric constant is applied to improve the accuracy of the ToA estimates. For example, the effective dielectric constant for the scene shown in Fig. 14 is approximated based on the ToA from each anchor to known reference locations (R1 and R2). For each anchor i, the ToA to the known reference location j is calculated from full-wave simulations (call it $t_f^{i,j}$). Then, the dielectric constant is computed as $\varepsilon_r^{i,j} = [t_f^{i,j}/t_e^{i,j}]^2$, where $t_e^{i,j}$ is the exact ToA computed from the known locations of the anchors and reference locations. The effective dielectric constant for the whole scene is then approximated by taking the

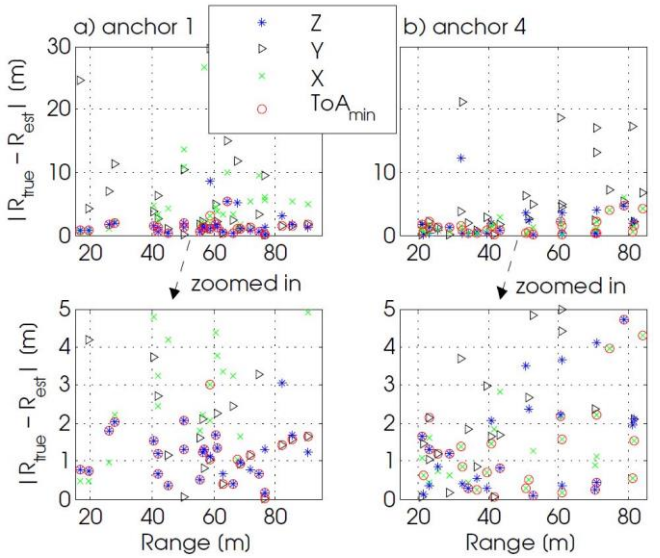


Fig. 16. The estimated and true location of the mobile node are plotted along with the anchor locations on the building map schematic.

average of the estimates from multiple reference locations (2 reference locations are used). The corrected ToA from anchor i to unknown location k is then computed as $t_c^{i,k} = t_f^{i,k}/[\varepsilon_r^i]^2$.

From the ToA estimates, the ranges from each anchor to the mobile node are computed. We then estimate the location of the mobile node using a weighted squared range least squares (WSR-LS) approach which provides an exact solution for the coordinate vector of the mobile node from the range estimates. It has been shown through numerical simulations that the WSR-LS approach is more accurate than existing approximations such as the unconstrained squared range-based LS estimate (USR-LS).

We investigate the proposed approach via a full-wave simulation of a realistic indoor/outdoor scenario. The scene shown in Fig. 14 consists of several buildings where both LoS and NLoS channels are considered. Concrete building walls modeled as lossy dielectric layers and steel reinforcement are included in the simulations. The antenna height on the mobile and reference nodes is 0.5m and the antenna height for the anchors is 7.5m. The simulations are performed using a finite difference time domain (FDTD) solver at center frequencies of 20MHz and 40MHz. The results shown use 5 anchors and 2 reference locations to estimate 29 unknown locations. Fig. 15 and 16 show the absolute value of the ranging error and a 2D map of the location estimates. The mean error without correcting for penetration delays with the effective dielectric constant is 1.23m. The mean error for the 29 locations of the mobile node goes down to 0.89m when the correction is applied. It should be noted that the FDTD solver introduces some error due to dispersion. The range error comparison among various antenna orientations given in Fig. 15 shows the advantage of polarization diversity.

IV. CONCLUSION AND FUTURE WORK

An electrically-small non-Foster matched antenna at low VHF is designed and tested. For the matching, a single NIC circuit is applied to a baseline passive antenna design in order to enhance its bandwidth and overall power efficiency. Stability analysis of the NIC with a suitably designed impedance model of the baseline antenna was conducted through circuit simulation. The antenna with the non-Foster matching was fabricated and its performance was measured and compared with that of the baseline design. The results show 3 dB power bandwidth improved by more than a factor of two with a power efficiency factor ranging from 1.3 to 5.9 within the bandwidth. The non-Foster matching design achieves an enhanced bandwidth and improved gain, exceeding the fundamental limit for passive antennas in the same form factor. This design will be utilized for low-power, short-range, low-VHF applications.

We have also presented an extremely simple method for localization, requiring a minimum of sensors and no calibration, synchronization, offline training, or array processing. By coupling the benign propagation of lower VHF with rigorous statistical modeling to handle correlations and outliers, we obtain useful point and uncertainty estimates of location. Our approach shows promise as a rapidly-deployable method for tracking and localization in complex scenarios. Further directions include optimizing beacon placement and motion of the mobile agent to improve accuracy.

The potential of utilizing polarization diversity to improve ToA based ranging and localization in harsh propagation scenarios leveraging the low frequency near-ground channel is investigated via full-wave simulations. The proposed technique exploits spatial and polarization diversity via multiple anchors co-located polarization diversity antennas. By applying a penetration delay compensation via an empirically derived effective dielectric constant, we show that the mean localization error can be reduced to 0.89m, a 28% improvement in accuracy compared to the result without correction. We are currently planning experiments to validate the proposed approach.

REFERENCES

[1] S. E. Sussman-Fort and R. M. Rudish, "Non-Foster impedance matching of electrically-small antennas," *IEEE Trans. Antennas Propag.*, vol. 57, no. 8, pp. 2230–2241, Aug. 2009.

[2] M. M. Jacob and D. F. Sievenpiper, "Gain and Noise Analysis of Non-Foster Matched Antennas," *IEEE Trans. Antennas Propag.*, vol. 64, no. 12, pp. 4993-5004, Dec. 2016.

[3] C. R. White, J. S. Colburn, R. G. Nagele, "A non-Foster VHF monopole antenna," *IEEE Antennas Wireless Propag. Lett.*, vol. 11, pp. 584-587, Jun. 2012.

[4] J. Choi, F. T. Dagefu, B. M. Sadler, and K. Sarabandi, "Electrically Small Folded Dipole Antenna for HF and Low-VHF Bands," *IEEE Antennas Wireless Propag. Lett.*, vol. 15, pp. 718-721, Aug. 2015.

[5] J. Wang, R. Ghosh, and S. K. Das, "A survey on sensor localization", *Journal of Control Theory and Applications*, vol. 8, pp. 2-11, 2010.

[6] G. Verma, F. Dagefu, B. M. Sadler, and K. Sarabandi, "Direction of arrival estimation with the received signal strength gradient at the lower VHF band," *IEEE International Symposium on Antennas and Propagation*, pp. 1695-1696, 2016.

[7] F. T. Dagefu, G. Verma, C. R. Rao, P. Y. Yu, J. R. Fink, B. M. Sadler, and K. Sarabandi, "Short Range Low VHF Channel Characterization in cluttered environments," *IEEE Trans. Ant. & Prop.*, vol. 63, no. 6, 2015.

Metrics

Awards and/or Recognitions

Honorable mention award, IEEE AP-S Student Paper Competition, 2016.

Refereed/Journal Papers

F. Dagefu, G. Verma, J. Choi, B. Sadler, and K. Sarabandi, "Low-Power Low-Frequency Communications in Complex Environments with Miniature Antennas," *IEEE Antennas and propagation Magazine*, to be submitted, 2017.

J. Choi, F. Dagefu, B. Sadler, and K. Sarabandi, "ZigBee-Based Low-VHF Enhanced Networking in Complex Environments," *IEEE Access, under review.*

J. Choi, F. Dagefu, B. Sadler, and K. Sarabandi, "A Miniature Non-Foster Matched Antenna for Low-Power Communications," *IEEE Antennas and Wireless Propagation Letters*, to be submitted, 2017.

K. Sarabandi, J. Choi, A. Sabet, and K. Sabet, "Pattern and Gain Characterization Using Non-Intrusive Very-Near-Field Electro-Optical Measurements over Arbitrary Closed Surfaces," *IEEE Transactions on Antennas and Propagation*, vol. 65, pp-489-497, February 2017.

J. Choi, F. Dagefu, B. Sadler, and K. Sarabandi, "Electrically Small Folded Dipole Antenna for HF and Low-VHF Bands," *IEEE Antennas and Wireless Propagation Letters*, vol. 15, pp. 718-721, August 2015.

F. Dagefu, J. Choi, M. Sheikhsofla, B. Sadler, and K. Sarabandi, " Performance Assessment of Lower VHF Band for Short Range Communication and Geolocation Applications", *Radio Science*, vol. 50, pp.443-452, May 2015.

F. Dagefu, G. Verma, C. Rao, P. Yu, B. Sadler, K. Sarabandi, "Short-Range Low-VHF Channel Characterization in Cluttered Environments", *IEEE Transactions on Antennas and Propagation*, vol. 63, no. 6, pp. 2719-2727, June 2015.

Conference Presentations/proceedings

J. Choi, F. T. Dagefu, B. M. Sadler, and K. Sarabandi, "A Non-Foster Matched Dipole for A Low-VHF Mobile Transmitter System," *2017 IEEE AP-S/URSI*, San Diego, CA, July 2017.

G. Verma, F. T. Dagefu, B. M. Sadler, and K. Sarabandi, "Localization via the Received Signal Strength Gradient at Lower VHF," *2017 IEEE AP-S/URSI*, San Diego, CA, July 2017.

F. T. Dagefu, G. Verma, Richard Kozick, and K. Sarabandi, "Full-wave Analysis of Time of Arrival Based Localization with Polarization Diversity," *2017 IEEE AP-S/URSI*, San Diego, CA, July 2017.

J. Choi and K. Sarabandi, "HF/VHF Antenna Characterization from Very-Near-Field Measurements over Arbitrary Closed Surfaces," *USNC-URSI National Radio Science Meeting*, Boulder, CO, USA, January 2016.

J. Choi and F. T. Dagefu, B. M. Sadler, and K. Sarabandi, "A Compact, Low-Power, Low-VHF Radio for Mobile and Wireless Communication Applications," *2016 IEEE AP-S/URSI*, Fajardo, Puerto Rico, June 2016.

J. Choi and K. Sarabandi, "Non-Foster Impedance Matching of an Extremely Small Folded-Dipole Antenna at Low-VHF Band," *2016 IEEE AP-S/URSI*, Fajardo, Puerto Rico, June 2016.

J. Choi and K. Sarabandi, "Full-Spherical Radiation Pattern Evaluation of Low Frequency Antennas Using a Novel Very-Near-Field Electro-Optical System," *2016 IEEE AP-S/URSI,* Fajardo, Puerto Rico, June 2016.

F. T. Dagefu, G. Verma, J. Choi, B. M. Sadler, and K. Sarabandi, "Empirical and Simulation based Analysis of Propagation Mechanisms at Low Frequencies," *2016 IEEE AP-S/URSI,* Fajardo, Puerto Rico, June 2016.

G. Verma, F. Dagefu, B.M. Sadler, K. Sarabandi, "Direction of arrival estimation with the received signal strength gradient at the lower VHF band," *2016 IEEE AP-S/URSI,* Fajardo, Puerto Rico, June 2016.

Joint Papers with ARL presented or accepted– duplication with journals/ conferences expected.

G. Verma, F. T. Dagefu, B. M. Sadler, and K. Sarabandi, "Localization via the Received Signal Strength Gradient at Lower VHF," *2017 IEEE AP-S/URSI*, San Diego, CA, July 2017.

F. T. Dagefu, G. Verma, Richard Kozick, and K. Sarabandi, "Full-wave Analysis of Time of Arrival Based Localization with Polarization Diversity," *2017 IEEE AP-S/URSI,* San Diego, CA, July 2017.

F. T. Dagefu, G. Verma, J. Choi, B. M. Sadler, and K. Sarabandi, "Empirical and Simulation based Analysis of Propagation Mechanisms at Low Frequencies," *2016 IEEE AP-S/URSI,* Fajardo, Puerto Rico, June 2016.

G. Verma, F. Dagefu, B.M. Sadler, K. Sarabandi, "Direction of arrival estimation with the received signal strength gradient at the lower VHF band," *2016 IEEE AP-S/URSI,* Fajardo, Puerto Rico, June 2016.

Invention Disclosures
Title

Patent Applications Filed
Title

Patents Awarded
Title

PhDs Award
Name(s): Jihun Choi

Masters Awarded
Name(s) Morteza Sheikhilsofla

Students supported fully or partially (by name)
MS: Lisa Smith
PhD: *Jihun Choi*
Post Docs: *Name(s)*
Master Degree: *Name(s)*
PhDs Awarded:

MAST CTA Member Rotations into ARL
Name(s)

ARL Researchers Working at Member Sites
Name(s)

Workshops and Seminars Conducted
Title

MAST Quarterly Metrics Summary Table (Numbers in this table should correspond to the list above):

MAST FINAL REPORT METRICS TABLE	
Awards (Miscellaneous Recognitions and Awards)	1
Refereed/Journal papers	7
Conference presentations/proceedings	9
Joint Papers with ARL	4
Invention disclosures	0
Patent applications filed	0
Patents awarded	0
PhDs Award	1
Masters Awarded	1
Post Docs Supported	0
Master Degree Students Supported	2
Undergraduate Students Supported	0
CTA Member rotations into ARL	0
Visiting Researchers from ARL to the Member Sites	0
Employees hired by ARL from CTA students supported	0

Short Range HF Communication in Highly Cluttered Environments

Task: CNC-17-1.1

Kamal Sarabandi, Jihun Choi
Dept. of Electrical and Computer Engineering
The University of Michigan
Ann Arbor, USA
saraband@umich.edu

Brian M. Sadler, Fikadu Dagefu
Army Research Laboratory
Adelphi, USA
brian.m.sadler6.civ@mail.mil

***Abstract*—We demonstrate bandwidth enhancement of a highly miniaturized low-VHF antenna via a non-Foster matching technique for low-power, short-range, wireless transmission systems. A recently developed miniature antenna ($0.013\,\lambda_0 \times 0.013\lambda_0 \times 0.02\lambda_0$ at 40 MHz) for this application is shown to have much higher radiation efficiency compared to similar-sized antennas, but its bandwidth is rather narrow. In order to improve the bandwidth while maintaining the best possible efficiency for the given size of the passive antenna, we introduces a non-Foster matching network realized by a negative impedance converter (NIC). We also presents a trade-off study between degree of bandwidth enhancement and the stability of the NIC circuit. Having a power efficiency advantage, a fabricated non-Foster matched antenna achieves a fractional bandwidth of about 2.15 %, an improvement of more than twice the bandwidth of the passive version. We also propose a minimalistic design for node localization based on a small number of access points. Our method uses the received signal strength (RSS) gradient as a direction of arrival (DoA) estimator. Though non-uniform shadowing and multipath adversely affect RSS, these are much less at lower VHF. We show through high fidelity simulations that our approach has promise as a rapidly-deployable localization system. Furthermore, we investigate a time of arrival based localization technique for nodes deployed in complex GPS-challenged environments. The proposed technique exploits spatial and polarization diversity via anchors equipped with polarization diversity antennas operating at low frequencies. Empirically derived effective dielectric constants are used for penetration delay compensation. Our full-wave simulations of a complex scene suggest that sub-meter localization error can be achieved, with low-power transmission at upper HF and lower VHF bands.**

***Index Terms*—Bandwidth Enhancement, HF band, low-VHF Band, localization, non-foster matching technique, polarization diversity, received signal strength, small antennas**

I. Introduction

Reliable communication is of paramount importance in many military applications. In scenarios where the communication environment is complex, such as urban or indoor environment, signal attenuation and multipath poses stringent conditions on power and signal processing. This situation is encountered by MAST small robotic platforms in all three scenarios under consideration including: 1) Urban search and mapping, 2) Cave or demolished building search and mapping, and 3) Perimeter defense. Radio communication is far more complex for such small robotic platforms since important wireless infrastructure, such as base stations with predetermined proper positions, are absent and high transmit power is not available. Much effort has been devoted over many decades to design modulation schemes appropriate for different communication bands to overcome challenges introduced by the communication channel and factors such as noise and interference. In general it is understood that channel complexity is far less at lower frequencies since significant signal penetration through/over obstacles is possible and reflection, scattering, and diffraction is far lower than what is encountered at upper UHF and higher bands. High penetration through soil and building materials allows for direct communication through caves and multiple layers of walls and buildings. Reduced reflection, scattering, and diffraction will be responsible for eliminating severe signal fluctuations, and fading which usually put stringent requirements on power, system bandwidth, as well as waveform and modulation. The severe signal fluctuations and intermittent signal drop-off that are unavoidable at higher frequencies are significantly reduced at these lower frequencies due their reduced reflection and scattering. It should be mentioned, however, that multipath effect is the main source of connectivity between a transmitter and a receiver in a complex environment at higher frequencies, where line-of-sight signal is extremely weak. On the other hand, at lower frequencies, the size of the conventional antennas is very large, making application of such systems limited to stationary nodes or, in the case of mobile nodes, the transmit power has to be high enough to compensate for the poor efficiency of the antennas. With advances in the design of small antennas, a careful analysis of propagation and scattering at the HF range using accurate physics based simulations as well as measurements in realistic scenarios is vital. Based on this analysis, the potential of using the upper HF band for applications in HF *ad hoc* networks can be assessed. Such a system can provide long distance coverage unhampered by walls, buildings, and foliage albeit being limited by instantaneous bandwidth. As mentioned before the limited bandwidth may not be a serious issue in the absence of fading.

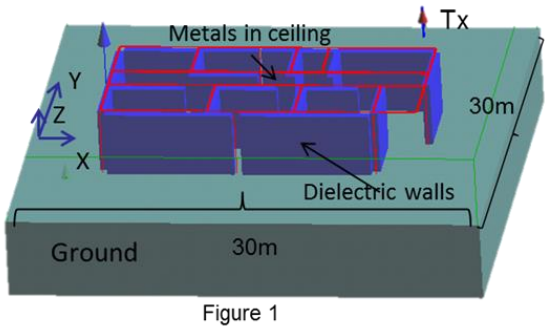


Fig. 1. Example of a single story office building with metallic frame.

With proper design of miniaturized antennas, it is possible to achieve multi-band antennas without significantly increasing their size.

II. RESEARCH ISSUE AND TECHNICAL APPROACH

A. Research Issue

Physics-based indoor propagation analysis in the HF band complemented by measurement results in complex indoor scenarios as well as the development of extremely small, light-weight HF antennas having reasonable gain and multi-band capability will play a vital role in supporting reliable direct communication among MAST platforms in complex environments including caves, underground, and through multiple layers of walls or buildings. Combining such communication nodes to perform beam-forming in an *ad hoc* manner among distributed nodes of MAST platforms allows very low power operation in very difficult propagation environments. The latter will be accomplished in collaboration with algorithms developed by Prof. Madhow of UCSB and Prof. Raghu Mudumbai of U. Iowa.

B. Technical Approach

As mentioned earlier, to study the pros and cons of an *ad hoc* wireless system at upper HF band, we will need to study the wave propagation behavior in complex propagation environments such as indoor and urban settings. In such environments building walls, ceiling and floors are the dominant scatterers. Fig. 1 shows a model of a complex building structure made from concrete cinder blocks where a transmitter is positioned outside the building (ceiling not shown). The ground is modeled as a half-space dielectric medium while the walls are modeled as thin dielectric slabs. A concrete ceiling is also included in the model. Realistic dielectric values are used for the walls, ground and ceiling. To make the environment more complex, a metallic frame is considered as support for the dielectric walls and ceiling. A short dipole located outside is utilized as a Tx antenna and field coverage inside and outside the building is computed as a function of frequency by utilizing accurate solutions. For such building walls, electromagnetic waves at high frequencies (UHF band and higher) will experience significant scattering. Basically, each ray that hits the walls will scatter in multiple directions. Each of the reflected

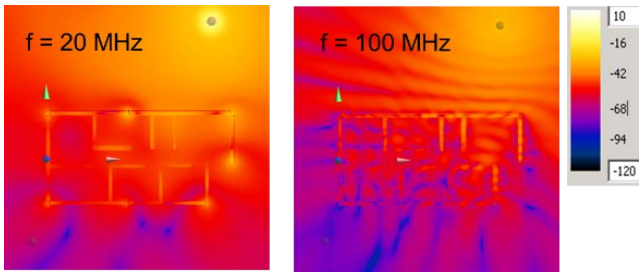


Fig. 2. Field coverage on a horizontal plane which is 1.5 m (=λ/10) above the ground. The left and right figures show, respectively, the field coverage at 20 MHz and 100 MHz, making communications at such low frequencies more reliable.

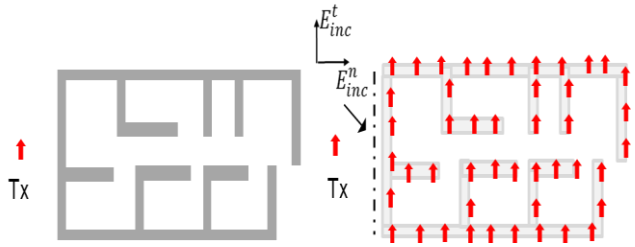
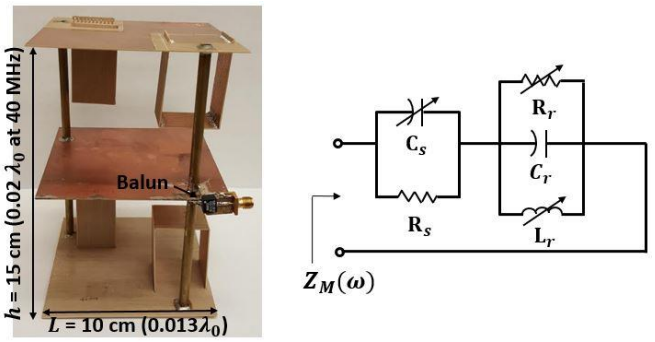


Fig. 3. The Proposed HF propagation model utilizes Rayleigh-Gans approximations in conjunction with a near-ground propagation model to take into account scattering from building walls. Polarization currents are used to capture the effect of the walls.

rays will bounce off of a different wall and multiply into many other rays, creating a highly fluctuating field distribution within the building. At lower frequencies (VHF and lower) models based on ray tracing do not work. For small buildings full-wave analysis methods such as the finite difference time domain (FDTD) method can be implemented on high performance computers. The signal coverage on a plane parallel to the ground at a height z_o = 1.5m is plotted for the building (see Fig. 1) based on a FDTD full-wave simulation for two different frequencies (20MHz and 100MHz) as shown in Fig. 2.

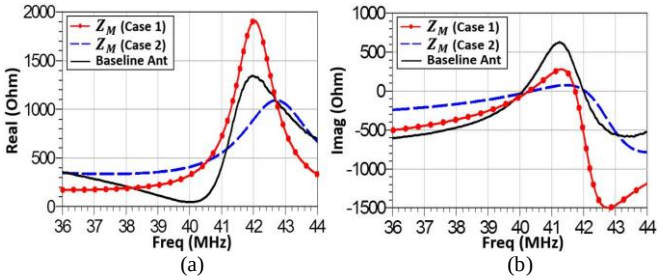
As expected, when a vertical dipole is used as Tx, the effect of the vertical metallic cylinders (found in walls in some non-residential buildings) can be seen in the field coverage at both frequencies. However, the horizontal metallic cylinders are not significant since the radiated field is vertically polarized. It is clear that even at 100MHz there is significant multipath and attenuation. At 20MHz, however, the walls appear as thin, effectively homogeneous, dielectric slabs with small reflectivity and high transmissivity. Under such conditions multi-path is minimized and the signal can propagate through walls with little attenuation. Another advantage is that the free-space path-loss, which increases with the square of frequency, is lower at lower frequencies. We also need to consider issues for near-ground propagation.

For transmitters and receivers near the ground the field decays as $1/R^2$ and the coefficient is proportional to the relative height (compared to the wavelength) of the transmitter and receiver. However, this degradation is not as important as the other two factors. We will examine different propagation scenarios for estimation of path loss using full-wave solutions for medium-size indoor scenarios and develop an efficient HF indoor and outdoor propagation model to analyze relatively large indoor and outdoor settings.



(a) (b)

Fig. 4. Baseline miniature passive antenna and its impedance model $Z_M(\omega)$.



(a) (b)

Fig. 5. Impedance model examples for the baseline antenna: (a) real and (b) imaginary parts of $Z_M(\omega)$ together with measured response of the baseline antenna.

The incident fields on the walls that have line-of-sight (LoS) with the transmit (Tx) antenna will be calculated by taking into account the near-ground effects. The fields will then be converted to equivalent currents using the Rayleigh or Rayleigh-Gans formulations, together with the current on the Tx antenna, will be used to illuminate the other non-line-of-sight (NLoS) walls, and calculate the field coverage inside and outside the building structure (see Fig. 3). Once the fields in the building are calculated, we will use equivalent currents and the analytical wave approach to calculate the field outside and over long distances. Parameters of interest include attenuation and field coverage, field fluctuation statistics, and coherence bandwidth for multiple MAST scenarios.

We will also investigate issues related to design of miniaturized antennas with enhanced bandwidth for operation at upper HF frequencies. Recently, we have designed and demonstrated performance of an extremely small antenna with dimensions as small as λ/100 by λ /100 by λ /300. The proposed antenna was not only small in size but also has very small weight (~15 g) since it was designed without any dielectric substrate unlike many of the state-of-the-art antennas. This antenna is capable of radiating purely vertical polarization with a relatively high gain of -18dBi. Investigations will be carried out to make improvements on this design, by increasing its bandwidth and efficiency. Tradeoff studies related to the center frequency, antenna size, data rate, and power consumption will also be carried out.

To tackle the challenge associated with the relatively small bandwidth of extremely small antennas such as the one described above, the use of multi-band antennas will be pursued. One possible approach is to design an array of a few of the miniaturized antennas operating at HF frequencies close to each other (e.g. 20 MHz, 24MHz, 28MHz) with one of the design goals being minimizing the overall size of the array. By utilizing the center frequencies as well as the first few harmonics, the available frequency band could be increased dramatically. The use of closely positioned narrowband antennas to increase bandwidth by taking advantage of mutual coupling between antenna elements is another approach that will be investigated.

III. RESEARCH ACCOMPLISHMENTS

A. *Bandwidth Enhancement of the Miniature Low-VHF Antenna*

In [1]-[3] the non-Foster matching technique is applied to unmatched small monopole or loop antennas. This technique utilizes a single negative capacitor and/or a single negative inductor produced by a negative impedance converter (NIC) to cancel out the reactive components of the antennas over a wide frequency range. However, for our baseline antenna (Fig. 4(a)), reactance cancellation is more challenging due to the steep impedance variation near the resonant frequency. Our goal is to broaden the effective bandwidth of the baseline passive antenna design while providing the best possible radiation efficiency for a given form factor. An advantage of this approach is that a single NIC can be applied directly to the antenna Balun, enhancing the overall power efficiency compared to using two NICs as in [1]. For its successful realization, challenging issues in antenna impedance modeling and stability with respect to system sensitivity are involved, which control the degree of bandwidth enhancement.

The extent of bandwidth enhancement by means of the proposed non-Foster matching technique critically depends on the accurate representation of the antenna impedance model. This will allow design of a NIC circuit to ameliorating the impedance variation of the antenna around its resonance. The baseline passive antenna has relatively narrow bandwidth (the 3 dB bandwidth of this design is 0.35 MHz or 0.87 %), which indicates a strong dependence of input impedance on frequency. Usually antenna impedance near the resonance is simply modeled by an RLC circuit. To improve the bandwidth by a significant amount a simple RLC circuit model is not effective. A more appropriate model of the antenna at hand is shown in Fig. 4(b), where impedance model $Z_M(\omega)$ consists of two series parallel circuits and can be expressed as

$$Z_M(\omega) = R_r \parallel j\omega L_r \parallel \frac{1}{j\omega C_r} + R_s \parallel \frac{1}{j\omega C_s}, \qquad (1)$$

where ω is the angular frequency. Fig. 5 illustrates examples of the real and imaginary parts of $Z_M(\omega)$, together with the measured antenna impedance. Case 1 parameters are $C_S = 6\,pF, R_s = 3.5\,k\Omega,\ R_r = 1.8\,k\Omega, L_r = 239\,nH,$ and $C_r =$

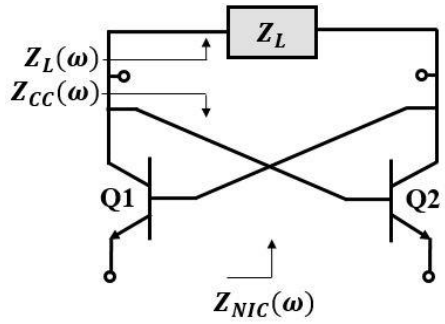


Fig. 6. A schematic of a cross-coupled pair consisting of two NPN bipolar junction transistors (BJT). $Z_L(\omega)$ is the impedance to be inverted with its negative value.

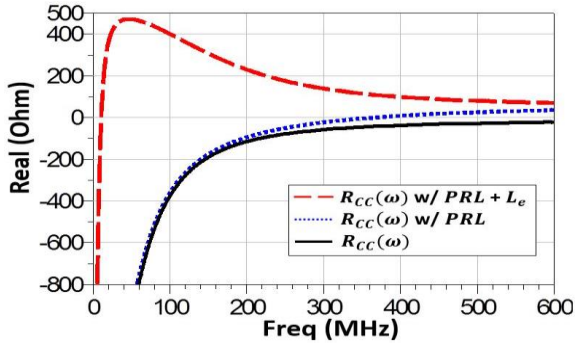


Fig. 7. The real part of the cross-coupled pair in Fig. 6 is shown, which is dependent on a parallel inductor-resistor *PRL* along the cross-coupled path and an inductor L_e in series with a resistor at emitters of the BJTs.

$60\ pF$. Case 2 parameters are $C_S = 10\ pF, R_s = 810\ \Omega, R_r = 880\ \Omega, L_r = 204\ nH$, and $C_r = 68\ pF$.

Note that the modeling parameters affect both the degree of bandwidth enhancement and system stability. For example, comparing the imaginary part of $Z_M(\omega)$ for the two cases shows that Case 1 better represents the baseline antenna. Therefore, an NIC circuit can be designed for better reactance cancellation through the NIC, but its faster resistance variation with frequency may lead to a stability issue with the non-Foster matching network (this will be discussed in the following section). This illustrates that the modeling parameters must be optimally chosen to achieve the maximal bandwidth under stable performance of the system.

A Bipolar Junction Transistor (BJT) cross-coupled pair (see Fig. 6) is widely used to produce a negative resistance in oscillator design or to function as a NIC. However, a careful stability analysis of the NIC is needed for non-Foster matching applications, due to the internal positive feedback that can lead to circuit oscillation. We design the matching circuit with the NIC to meet the stability criterion described below, simulate the circuit, and apply the Nyquist stability test to verify the validity of the design. The first step in this process is to characterize the impedance of the cross-coupled pair $Z_{CC}(\omega) = R_{CC}(\omega) + jX_{CC}(\omega)$ with a DC biasing circuit as a function of frequency in the small-signal regime. Since the cross-coupled pair itself generates a negative resistance, $R_{CC}(\omega) < 0$ when a passive load impedance $Z_L(\omega) = R_L(\omega) + jX_L(\omega) = \infty$ is an open

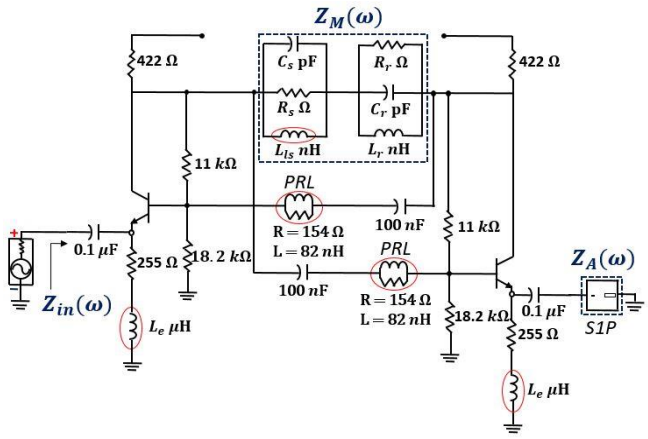


Fig. 8. A schematic of the non-Foster matching network designed to ensure stable system performance.

circuit (this is plotted as a black dashed line in Fig. 7). Also, to avoid oscillation conditions when $Z_L(\omega)$ is connected, the net resistance of the network must be positive, expressed as

$$Re[Z_{CC}(\omega) + Z_L(\omega)] > 0 \ or \ |R_{CC}(\omega)| < R_L(\omega). \qquad (2)$$

Here, $Z_L(\omega)$ corresponds to the aforementioned $Z_M(\omega)$ for non-Foster matching of the baseline antenna. In order to satisfy (2) at all frequencies, we proceed as follows.

To make $R_{CC}(\omega)$ positive at higher frequencies, the feedback factor along each cross-coupled path is adjusted by adding a parallel inductor-resistor *PRL* ($124\ \Omega \parallel 82\ nH$) in series with a feedback capacitor. Furthermore, an inductor L_e ($10\ \mu H$) in series with the emitter resistor is added for each BJT. With the added components, $R_{CC}(\omega)$ becomes positive above 12 MHz as depicted in Fig. 7. To fulfill (2) at lower frequencies, $Z_M(\omega)$ is modified by adding an inductor L_{ls} in parallel with R_s and C_s to produce high impedance below 12 MHz.

Fig. 8 illustrates a schematic of the designed matching network and the relevant components added to the design in Fig. 6 are circled. Here, $Z_{in}(\omega)$ represents the sum of the impedance of the floating NIC, $Z_{NIC}(\omega)$ (see Fig. 6), and the measured antenna impedance, $Z_A(\omega) = R_A(\omega) + jX_A(\omega)$. The total loop resistance must also be positive in order to avoid circuit oscillation. This design was simulated over a range of values for the elements in $Z_M(\omega)$, taking into account the component tolerances and reduction of the antenna impedance variation over the frequency range of interest while avoiding oscillations. Next, we conduct a large-signal S-parameter simulation to ensure overall robustness to nonlinear effects of the BJT with respect to large AC signals and variations in power supply values.

Based on the design described above, a non-Foster matching network was fabricated and characterized. In order to fine tune $Z_M(\omega)$ in the presence of antenna impedance variations arising from the measurement environment (e.g., proximity effects) and component and fabrication tolerances, mechanically adjustable elements are utilized (they are marked with arrows

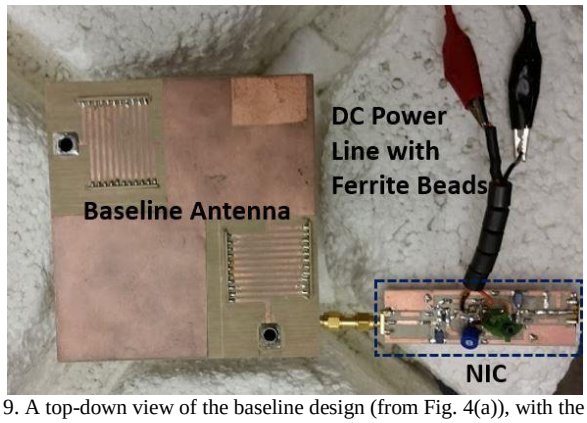


Fig. 9. A top-down view of the baseline design (from Fig. 4(a)), with the NIC connected directly to the antenna Balun.

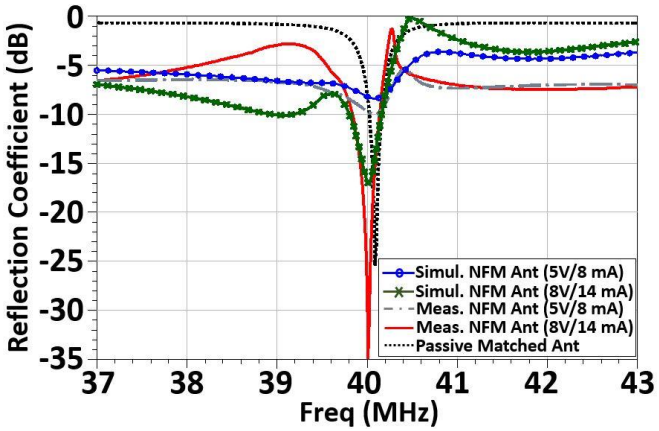


Fig. 10. Simulated and measured reflection coefficients as a function of power supply value expressed as voltage to current ratio.

in Fig. 4(b)). The fabricated NIC is connected directly to the antenna Balun along the null plane bisecting the antenna structure (see Fig. 4 and 9). This ensures that introducing the NIC does not affect the characteristics of the passive antenna such as its input impedance and radiation pattern [4]. Multiple ferrite beads are used on the DC power line for high frequency noise and parasitic suppression. They are also utilized on the coaxial feed cable to suppress common mode current flowing on the outer conductor of the coax.

The fabricated non-Foster matched antenna is characterized using a vector network analyzer. Large-signal measurements as a function of input power and power supply voltage are performed to examine nonlinear behaviors of the BJTs in the NIC. We first fix the input power to 7 dBm, corresponding to a relatively low power transmission in the low-VHF band. Keeping the input power fixed, we tested with different voltage/current ratios, measured reflection coefficients of the non-Foster matched antenna as a function of frequency, and compared them with circuit simulation results as shown in Fig. 10.

We also measured the transmission coefficient for the same cases as in Fig. 10. For this measurement, a dipole with a length of $0.2\,\lambda_0$ at 40 MHz is used as a receiving antenna. The transmission coefficients as a function of frequency are

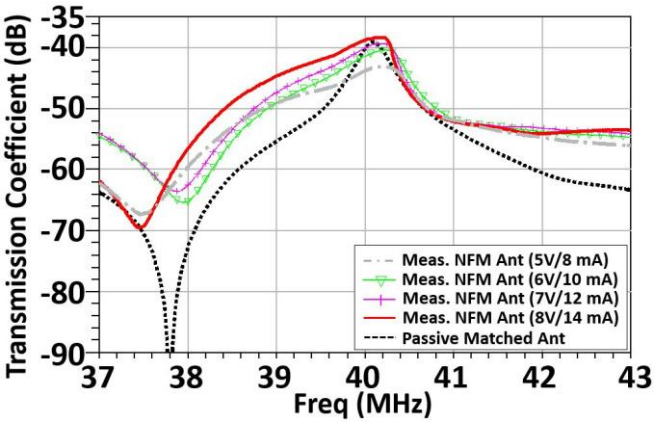


Fig. 11. Measured transmission coefficients as a function of power supply value. A $\lambda_0/5$ dipole is used as a receiving antenna.

TABLE I
COMPARISON OF ANTENNA MATCHING PERFORMANCE

Matching Type	3 dB BW	η_f (within the BW)
Passive	0.87 %	·
Non-Foster	2.15 %	1.3 to 5.9

depicted in Fig. 11.

The measurement results show the dependence of the impedance matching and DC bias. Note that, when the input AC signal is small compared to the DC bias, the cross-coupled pair properly functions as a NIC in the small-signal regime where balanced operation of the BJTs in the pair is maintained. This is why the matching performance is better for higher DC bias values as can be seen in the results. However, over the DC bias range indicated in Fig. 10, the cross-coupled pair turns into a regenerative circuit that gives rise to matching degradation and system instability. A similar behavior was also observed by sweeping the input AC signal while keeping a fixed DC bias (not shown here). Thus, the maximal bandwidth of the matching network under stable operating conditions can be determined with a proper DC bias for a given input power.

Since the non-Foster matched antenna requires DC power to operate the NIC, the overall power efficiency should be considered to evaluate the effectiveness of the proposed approach against the baseline passive design. We define a power efficiency factor η_f as

$$\eta_f = \frac{|S_{21(AM)}|^2}{|S_{21(PM)}|^2} \cdot \frac{P_{DC(PM)}}{P_{DC(AM)}}. \quad (3)$$

Here, $S_{21(AM)}$ and $S_{21(PM)}$ are measured transmission coefficients with active and passive matching, respectively. $P_{DC(PM)}$ is the DC power for transmission in the passive case, and for the active case $P_{DC(AM)}$ is the DC power for transmission plus the power for the active matching circuit. Aligning with low-VHF radio experiments, P_{DC} for the transmitter is set to be 3.3 V/ 240 mA, and P_{DC} for the NIC is

chosen to be 8 V/ 14 mA. With these values, at a 40 MHz center frequency, the active matching enhances the 3 dB bandwidth from 0.87 % to 2.15 % with a power efficiency advantage such that $\eta_f > 1$ over the entire bandwidth. The performance comparison between the passive and active matching is summarized in Table I.

B. Localization via the Received Signal Strength Gradient

Localization and tracking of targets are critical objectives in a variety of fields, e.g., in radar and mobile robotic systems. Use of global positioning system (GPS) is not always feasible due to cost and/or poor penetration. Many localization techniques rely on time of arrival (ToA), angle of arrival (AoA), or received signal strength (RSS). ToA-based methods based on time-of-flight of a wireless signal suffer from errors due to clock drift and variable processing delays. Classical AoA-based approaches rely on array-processing, requiring multiple synchronized antenna elements per node, and have many challenges, e.g., the location estimate is very sensitive to multipath. Classical RSS-based localization approaches attempt to estimate distance-to-source from the received power; this requires estimating various nuisance parameters (e.g., path loss exponent) and in complex scenes, RSS is corrupted by multipath and shadowing [5]. In this work, we take a fundamentally different approach; we use RSS to compute the direction of arrival (DoA) to fixed-location nodes (beacons), and fuse the DoA's to localize the node. This method relies on the basic fact that under ideal propagation, the RSS spatial gradient points towards the source. At microwave frequencies, huge fluctuations in RSS occur over small spatial scales due to small-scale fading. However, these effects are much less prevalent at lower VHF. In addition, we can employ statistical methods to greatly reduce the effect of RSS outliers. Our framework yields point estimates of the DoA as well as uncertainty regions. RSS is easy to measure, does not require an array to process, and since the gradient is a relative measure (a slope), it requires no calibration. Estimation of the RSS spatial gradient does require multiple spatially separated RSS measurements.

Our model for DoA estimation from RSS data is described in [6]. The RSS surface is locally modelled as a linear function of sampling location, and the gradient of the model is a DoA estimate. Notationally,

$$y = X\beta_i + \varepsilon \tag{4}$$

where y_i is a $N \times 1$ vector of RSS values of beacon i at N sampling locations, X is a N × 2 matrix of those RSS sampling locations, and $\varepsilon = N(0, \sigma^2 I_N)$, where I_N is the N × N identity matrix. We assume the presence of $B \geq 2$ beacons with known fixed locations. Thus, a node seeking to localize itself measures the RSS with respect to each of the B beacons over X. We adopt a Bayesian model over 1 which enables us to incorporate prior information on σ^2 and to reason probabilistically about the

Algorithm 1 RSS gradient-based localization

for m=1 to M **do**
  for i=1 to B **do**
    $\boldsymbol{\beta_i} \leftarrow \boldsymbol{p}(\boldsymbol{\beta_i}|\boldsymbol{y}, \boldsymbol{X})$
    $\boldsymbol{n_i} \leftarrow \begin{pmatrix} \mathbf{0} & \mathbf{-1} \\ \mathbf{1} & \mathbf{0} \end{pmatrix} \frac{\boldsymbol{\beta_i}}{\|\boldsymbol{\beta_i}\|}$
    $\boldsymbol{\omega_i} \leftarrow \|\boldsymbol{\beta_i}\|$
  end for
  $\boldsymbol{l_m} \leftarrow \left(\sum_i \boldsymbol{\omega_i} \boldsymbol{n_i} \boldsymbol{n_i^T}\right)^{-1} ((\sum_i \boldsymbol{\omega_i} \boldsymbol{n_i} \boldsymbol{n_i^T} \boldsymbol{b_i})$
end for

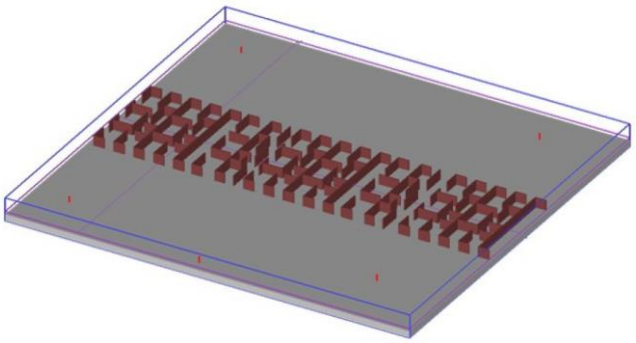


Fig. 12. FDTD simulation of an office building with realistic materials. All nodes are at assumed height of 0.1λ and at center frequency of 40 MHz. The total scene is 122 m × 115 m. Beacons (red dots) with known locations surround the building.

DoA estimate β_i, where i ranges from 1 to B. We can use previous studies characterizing the shadowing variance at lower VHF to derive distributions over σ^2. We use $p(y|X, \beta) = N(X\beta, \sigma^2 I_N)$, a Normal distribution. Using standard statistical theory, we can derive the posterior probability distribution $p(\beta_i|y, X)$ over each β_i.

We next combine $\{\beta_i\}_{i=1}^{i=B}$ into a localization estimate. In the general case, the node's orientation θ is unknown (with respect to the global coordinate system). θ is straightforwardly estimated; but for simplicity in the sequel we assume an estimate of it is known (e.g., via compass), so that the β_i are represented in the global coordinate system. Our localization algorithm is presented in Algorithm 1. For each of M Monte Carlo iterations, we sample all B direction vectors from their posterior distributions. We apply a weighted least squares procedure on these DoA's to obtain a point estimate of the node location l based on known beacon locations bi. M such Monte Carlo runs yield an empirical distribution over node location. We must account for correlations and outliers in the data to ensure the estimates are meaningful. Positive spatial correlation arises when nearby RSS samples encounter correlated shadowing. Without correction, the true uncertainty in the DoA (and hence localization) is underestimated. Outliers arise from fading and non-uniform shadowing, and bias the estimated DoA from the true source direction. We use standard statistical procedures for correcting spatial correlation and outliers; details will be described elsewhere.

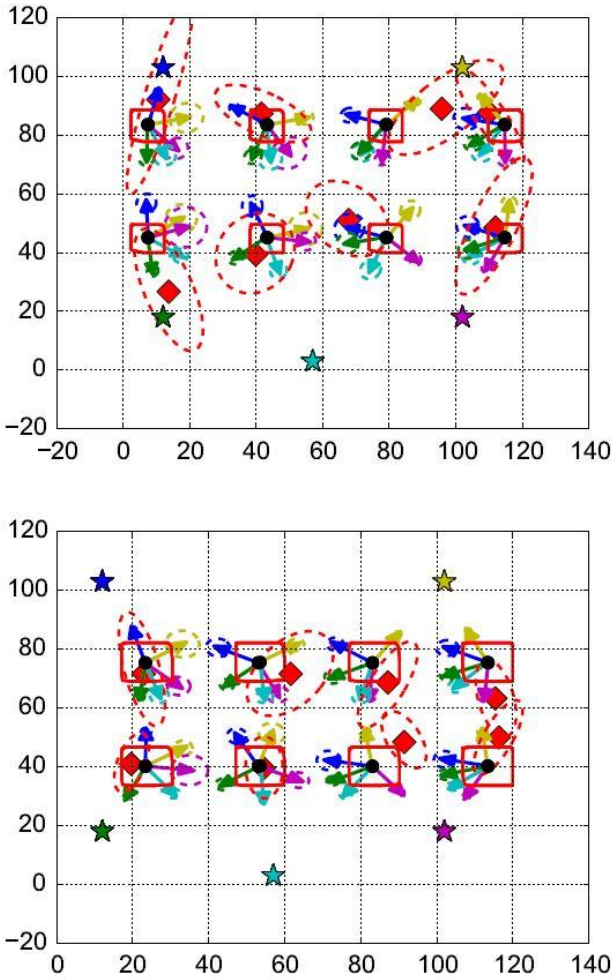


Fig. 13. Localization results for a mobile agent which samples 8 different sampling regions within the office scene in Fig. 12 of sizes (top) ≈ 10m × 10m and (bottom) ≈ 12m × 12m. Figure elements: beacon (colored star); estimated DoA based on RSS (arrow colored to match corresponding beacon); DoA uncertainty region (colored dashed ellipse encircling DoA arrow); convex hull of the sampling region (rectangular red region); true location (i.e., centroid of sampling region) (black dot); location estimate (red diamond); location uncertainty ellipse (dashed red ellipse). Top: Median localization error is 10:1 m. Bottom: Median localization error is 7:4 m.

We have conducted high-fidelity simulations at lower VHF using a finite-difference time-domain (FDTD) full wave solver in a complex indoor office environment shown in Fig. 12. Fig. 13 shows the results; the caption has further details. Note that as the RSS sampling region grows, the typical localization error decreases and the uncertainty ellipses shrink.

C. Full-Wave Analysis of Time of Arrival Based Localization with Polarization Diversity

Accurate localization of mobile nodes in GPS-challenged scenarios has various civilian and military applications including navigation of robotic platforms deployed in complex infrastructure-poor environments for situational awareness as well as search and rescue missions. Conventional techniques that operate at microwave frequency bands either have poor accuracies or require extensive training and/or site specific measurement database and often assume the availability of infrastructure. For near-ground nodes deployed in such environments, recent studies suggest that lower frequencies (20-120 MHz) offer significant advantages due to the presence of a dominant line-of-sight (LoS) like path even when there are a multitude of scatterers between the nodes. This is due to the improved penetration and reduced multipath levels [7]. The presence of a dominant direct path makes the near-ground low frequency channel a good candidate to improve the accuracy of time of arrival (ToA) and array based localization techniques.

However, the presence of electrically large conductive scatterers creates signal components that will introduce significant errors. The strength of the scattered components is a function of the electrical size, material characteristics, frequency, range and polarization. For example, a vertically oriented metallic cylinder will reflect a vertically polarized wave, but minimally affects horizontally polarized field, and vice versa. This suggests that, depending on the orientation of nearby scatterers, one polarization may be preferred to minimize scattering and preserve a dominant LoS-like channel. Another important factor is phase aberration due to penetration delays through obstacles, which can significantly bias the ToA based range estimate, resulting in a biased localization estimate. Here, we report our initial investigation of a ToA based localization technique that exploits polarization diversity and empirical penetration delay compensation.

We consider a ToA based localization technique that leverages the low frequency near-ground channel in conjunction with polarization diversity and penetration delay compensation to significantly reduce ranging and localization errors. A typical scenario of interest consists of a minimum of three anchors with known locations as well as a mobile node deployed in an indoor or urban scenario. We assume all nodes are equipped with co-located polarization diversity antennas and the anchors transmit known Gaussian pulses in all three antenna orientations. The mobile node captures the three components of the signal from the various anchors. We assume orthogonal transmission among the anchors, e.g. using code division multiple access. As alluded to in the previous section, although the near-ground low frequency channel has a LoS-like dominant component due to better penetration and reduced multipath [7], the presence of electrically large scatterers near the antennas can cause some polarizations to experience significant multipath depending on the orientation of the dominant obstacle.

For a given location of the mobile node, the received signal from each anchor will have three components corresponding to the three orientations of the antennas. The receiver employs a matched filter for each polarization, matched to the Gaussian pulse. Three different ToA estimates are then computed by taking the difference between receive pulse time for each polarization and reference transmit time. We then choose the smallest ToA as our estimate. This ensures that the polarization that experienced the least multipath is selected as the ToA estimate for that anchor.

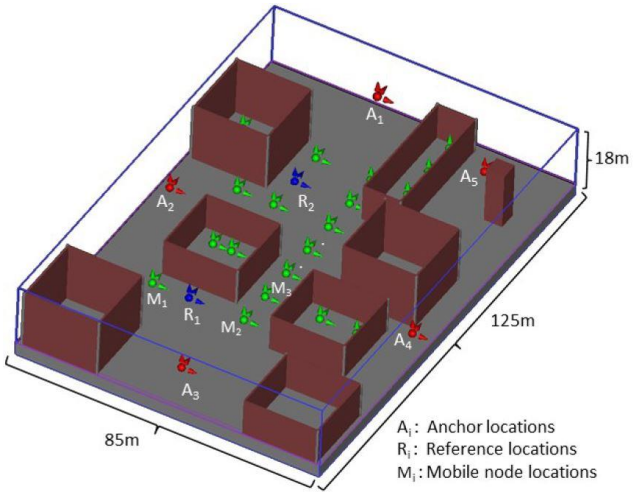


Fig. 14. An indoor-outdoor scenario consisting of dielectric and metallic scatterers with 5 anchors, 2 reference locations and 29 mobile node locations.

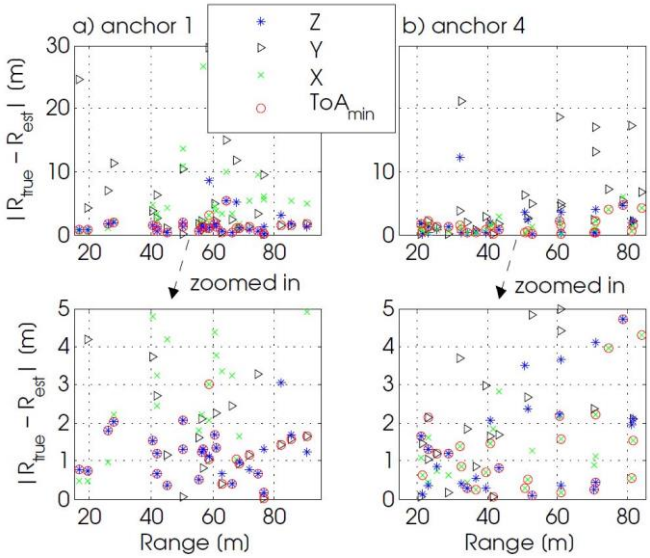


Fig. 15. The ranging errors based on ToA estimated from each polarization pairs are shown for anchors 1 and 4. The final estimate (denoted byToA_{min}) is computed by taking the smallest ToA estimate based on all 3 polarizations; note that the final choice varies with node location.

In the presence of multiple layers of large scatterers, even if there is a dominant direct path, the spatially varying penetration delay through multiple layers of scatterers could be significant. To combat this, an empirically estimated anchor-dependent effective dielectric constant is applied to improve the accuracy of the ToA estimates. For example, the effective dielectric constant for the scene shown in Fig. 14 is approximated based on the ToA from each anchor to known reference locations (R1 and R2). For each anchor i, the ToA to the known reference location j is calculated from full-wave simulations (call it $t_f^{i,j}$). Then, the dielectric constant is computed as $\varepsilon_r^{i,j} = [t_f^{i,j}/t_e^{i,j}]^2$, where $t_e^{i,j}$ is the exact ToA computed from the known locations of the anchors and reference locations. The effective dielectric constant for the whole scene is then approximated by taking the

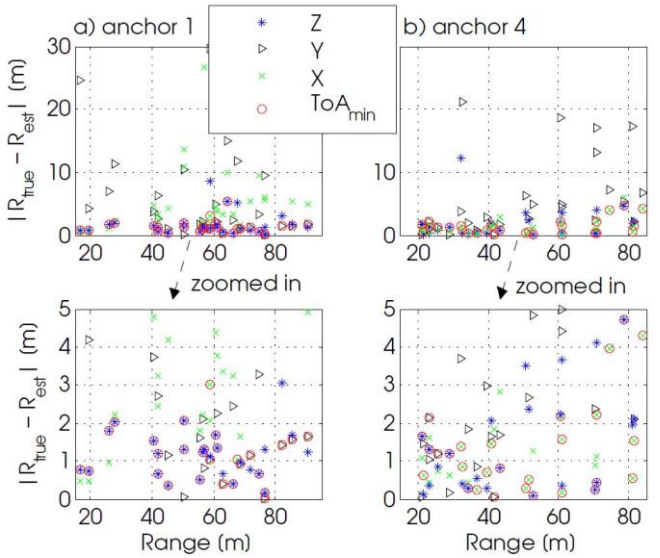


Fig. 16. The estimated and true location of the mobile node are plotted along with the anchor locations on the building map schematic.

average of the estimates from multiple reference locations (2 reference locations are used). The corrected ToA from anchor i to unknown location k is then computed as $t_c^{i,k} = t_f^{i,k}/[\varepsilon_r^i]^2$.

From the ToA estimates, the ranges from each anchor to the mobile node are computed. We then estimate the location of the mobile node using a weighted squared range least squares (WSR-LS) approach which provides an exact solution for the coordinate vector of the mobile node from the range estimates. It has been shown through numerical simulations that the WSR-LS approach is more accurate than existing approximations such as the unconstrained squared range-based LS estimate (USR-LS).

We investigate the proposed approach via a full-wave simulation of a realistic indoor/outdoor scenario. The scene shown in Fig. 14 consists of several buildings where both LoS and NLoS channels are considered. Concrete building walls modeled as lossy dielectric layers and steel reinforcement are included in the simulations. The antenna height on the mobile and reference nodes is 0.5m and the antenna height for the anchors is 7.5m. The simulations are performed using a finite difference time domain (FDTD) solver at center frequencies of 20MHz and 40MHz. The results shown use 5 anchors and 2 reference locations to estimate 29 unknown locations. Fig. 15 and 16 show the absolute value of the ranging error and a 2D map of the location estimates. The mean error without correcting for penetration delays with the effective dielectric constant is 1.23m. The mean error for the 29 locations of the mobile node goes down to 0.89m when the correction is applied. It should be noted that the FDTD solver introduces some error due to dispersion. The range error comparison among various antenna orientations given in Fig. 15 shows the advantage of polarization diversity.

IV. CONCLUSION AND FUTURE WORK

An electrically-small non-Foster matched antenna at low VHF is designed and tested. For the matching, a single NIC circuit is applied to a baseline passive antenna design in order to enhance its bandwidth and overall power efficiency. Stability analysis of the NIC with a suitably designed impedance model of the baseline antenna was conducted through circuit simulation. The antenna with the non-Foster matching was fabricated and its performance was measured and compared with that of the baseline design. The results show 3 dB power bandwidth improved by more than a factor of two with a power efficiency factor ranging from 1.3 to 5.9 within the bandwidth. The non-Foster matching design achieves an enhanced bandwidth and improved gain, exceeding the fundamental limit for passive antennas in the same form factor. This design will be utilized for low-power, short-range, low-VHF applications.

We have also presented an extremely simple method for localization, requiring a minimum of sensors and no calibration, synchronization, offline training, or array processing. By coupling the benign propagation of lower VHF with rigorous statistical modeling to handle correlations and outliers, we obtain useful point and uncertainty estimates of location. Our approach shows promise as a rapidly-deployable method for tracking and localization in complex scenarios. Further directions include optimizing beacon placement and motion of the mobile agent to improve accuracy.

The potential of utilizing polarization diversity to improve ToA based ranging and localization in harsh propagation scenarios leveraging the low frequency near-ground channel is investigated via full-wave simulations. The proposed technique exploits spatial and polarization diversity via multiple anchors co-located polarization diversity antennas. By applying a penetration delay compensation via an empirically derived effective dielectric constant, we show that the mean localization error can be reduced to 0.89m, a 28% improvement in accuracy compared to the result without correction. We are currently planning experiments to validate the proposed approach.

REFERENCES

[1] S. E. Sussman-Fort and R. M. Rudish, "Non-Foster impedance matching of electrically-small antennas," *IEEE Trans. Antennas Propag.*, vol. 57, no. 8, pp. 2230–2241, Aug. 2009.

[2] M. M. Jacob and D. F. Sievenpiper, "Gain and Noise Analysis of Non-Foster Matched Antennas," *IEEE Trans. Antennas Propag.*, vol. 64, no. 12, pp. 4993-5004, Dec. 2016.

[3] C. R. White, J. S. Colburn, R. G. Nagele, "A non-Foster VHF monopole antenna," *IEEE Antennas Wireless Propag. Lett.*, vol. 11, pp. 584-587, Jun. 2012.

[4] J. Choi, F. T. Dagefu, B. M. Sadler, and K. Sarabandi, "Electrically Small Folded Dipole Antenna for HF and Low-VHF Bands," *IEEE Antennas Wireless Propag. Lett.*, vol. 15, pp. 718-721, Aug. 2015.

[5] J. Wang, R. Ghosh, and S. K. Das, "A survey on sensor localization", *Journal of Control Theory and Applications*, vol. 8, pp. 2-11, 2010.

[6] G. Verma, F. Dagefu, B. M. Sadler, and K. Sarabandi, "Direction of arrival estimation with the received signal strength gradient at the lower VHF band," *IEEE International Symposium on Antennas and Propagation*, pp. 1695-1696, 2016.

[7] F. T. Dagefu, G. Verma, C. R. Rao, P. Y. Yu, J. R. Fink, B. M. Sadler, and K. Sarabandi, "Short Range Low VHF Channel Characterization in cluttered environments," *IEEE Trans. Ant. & Prop.*, vol. 63, no. 6, 2015.

Metrics

Awards and/or Recognitions

Honorable mention award, IEEE AP-S Student Paper Competition, 2016.

Refereed/Journal Papers

F. Dagefu, G. Verma, J. Choi, B. Sadler, and K. Sarabandi, "Low-Power Low-Frequency Communications in Complex Environments with Miniature Antennas," *IEEE Antennas and propagation Magazine*, to be submitted, 2017.

J. Choi, F. Dagefu, B. Sadler, and K. Sarabandi, "ZigBee-Based Low-VHF Enhanced Networking in Complex Environments," *IEEE Access, under review.*

J. Choi, F. Dagefu, B. Sadler, and K. Sarabandi, "A Miniature Non-Foster Matched Antenna for Low-Power Communications," *IEEE Antennas and Wireless Propagation Letters*, to be submitted, 2017.

K. Sarabandi, J. Choi, A. Sabet, and K. Sabet, "Pattern and Gain Characterization Using Non-Intrusive Very-Near-Field Electro-Optical Measurements over Arbitrary Closed Surfaces," *IEEE Transactions on Antennas and Propagation*, vol. 65, pp-489-497, February 2017.

J. Choi, F. Dagefu, B. Sadler, and K. Sarabandi, "Electrically Small Folded Dipole Antenna for HF and Low-VHF Bands," *IEEE Antennas and Wireless Propagation Letters*, vol. 15, pp. 718-721, August 2015.

F. Dagefu, J. Choi, M. Sheikhsofla, B. Sadler, and K. Sarabandi, " Performance Assessment of Lower VHF Band for Short Range Communication and Geolocation Applications", *Radio Science*, vol. 50, pp.443-452, May 2015.

F. Dagefu, G. Verma, C. Rao, P. Yu, B. Sadler, K. Sarabandi, "Short-Range Low-VHF Channel Characterization in Cluttered Environments", *IEEE Transactions on Antennas and Propagation*, vol. 63, no. 6, pp. 2719-2727, June 2015.

Conference Presentations/proceedings

J. Choi, F. T. Dagefu, B. M. Sadler, and K. Sarabandi, "A Non-Foster Matched Dipole for A Low-VHF Mobile Transmitter System," *2017 IEEE AP-S/URSI*, San Diego, CA, July 2017.

G. Verma, F. T. Dagefu, B. M. Sadler, and K. Sarabandi, "Localization via the Received Signal Strength Gradient at Lower VHF," *2017 IEEE AP-S/URSI*, San Diego, CA, July 2017.

F. T. Dagefu, G. Verma, Richard Kozick, and K. Sarabandi, "Full-wave Analysis of Time of Arrival Based Localization with Polarization Diversity," *2017 IEEE AP-S/URSI*, San Diego, CA, July 2017.

J. Choi and K. Sarabandi, "HF/VHF Antenna Characterization from Very-Near-Field Measurements over Arbitrary Closed Surfaces," *USNC-URSI National Radio Science Meeting*, Boulder, CO, USA, January 2016.

J. Choi and F. T. Dagefu, B. M. Sadler, and K. Sarabandi, "A Compact, Low-Power, Low-VHF Radio for Mobile and Wireless Communication Applications," *2016 IEEE AP-S/URSI*, Fajardo, Puerto Rico, June 2016.

J. Choi and K. Sarabandi, "Non-Foster Impedance Matching of an Extremely Small Folded-Dipole Antenna at Low-VHF Band," *2016 IEEE AP-S/URSI*, Fajardo, Puerto Rico, June 2016.

J. Choi and K. Sarabandi, "Full-Spherical Radiation Pattern Evaluation of Low Frequency Antennas Using a Novel Very-Near-Field Electro-Optical System," *2016 IEEE AP-S/URSI,* Fajardo, Puerto Rico, June 2016.

F. T. Dagefu, G. Verma, J. Choi, B. M. Sadler, and K. Sarabandi, "Empirical and Simulation based Analysis of Propagation Mechanisms at Low Frequencies," *2016 IEEE AP-S/URSI,* Fajardo, Puerto Rico, June 2016.

G. Verma, F. Dagefu, B.M. Sadler, K. Sarabandi, "Direction of arrival estimation with the received signal strength gradient at the lower VHF band," *2016 IEEE AP-S/URSI,* Fajardo, Puerto Rico, June 2016.

Joint Papers with ARL presented or accepted– duplication with journals/ conferences expected.

G. Verma, F. T. Dagefu, B. M. Sadler, and K. Sarabandi, "Localization via the Received Signal Strength Gradient at Lower VHF," *2017 IEEE AP-S/URSI,* San Diego, CA, July 2017.

F. T. Dagefu, G. Verma, Richard Kozick, and K. Sarabandi, "Full-wave Analysis of Time of Arrival Based Localization with Polarization Diversity," *2017 IEEE AP-S/URSI,* San Diego, CA, July 2017.

F. T. Dagefu, G. Verma, J. Choi, B. M. Sadler, and K. Sarabandi, "Empirical and Simulation based Analysis of Propagation Mechanisms at Low Frequencies," *2016 IEEE AP-S/URSI,* Fajardo, Puerto Rico, June 2016.

G. Verma, F. Dagefu, B.M. Sadler, K. Sarabandi, "Direction of arrival estimation with the received signal strength gradient at the lower VHF band," *2016 IEEE AP-S/URSI,* Fajardo, Puerto Rico, June 2016.

Invention Disclosures
Title

Patent Applications Filed
Title

Patents Awarded
Title

PhDs Award
Name(s): Jihun Choi

Masters Awarded
Name(s) Morteza Sheikhilsofla

Students supported fully or partially (by name)
MS: Lisa Smith
PhD: *Jihun Choi*
Post Docs: *Name(s)*
Master Degree: *Name(s)*
PhDs Awarded:

MAST CTA Member Rotations into ARL
Name(s)

ARL Researchers Working at Member Sites
Name(s)

Workshops and Seminars Conducted
Title

MAST Quarterly Metrics Summary Table (Numbers in this table should correspond to the list above):

MAST FINAL REPORT METRICS TABLE	
Awards (Miscellaneous Recognitions and Awards)	1
Refereed/Journal papers	7
Conference presentations/proceedings	9
Joint Papers with ARL	4
Invention disclosures	0
Patent applications filed	0
Patents awarded	0
PhDs Award	1
Masters Awarded	1
Post Docs Supported	0
Master Degree Students Supported	2
Undergraduate Students Supported	0
CTA Member rotations into ARL	0
Visiting Researchers from ARL to the Member Sites	0
Employees hired by ARL from CTA students supported	0

Sensor-based Motion Planning using Gaussian Mixtures and Rényi's Quadratic Entropy

CNC 16-17-2.1 Distributed Information Games for Autonomous Team Behavior

Shih-Ling Phuong, Nikolay Atanasov, Alejandro Ribeiro
Dept. of Electrical and Systems Engineering,
University of Pennsylvania,
Philadelphia, Pennsylvania,
{sphuong, atanasov, aribeiro}@seas.upenn.edu.

***Abstract*—This paper examines the problem of tracking a moving target by exploring an unknown environment using a mobile sensor. The sensor actively explores without mapping by modeling the state of the target as a Gaussian mixture model. This paper focuses on designing a control policy for the sensor which minimizes a utility function based on the uncertainty of the target's state. First, we present an information-theoretic utility function derived using Rényi's definition of entropy to quantify the uncertainty of the target's state given the sensor measurements. We then outline the genearl approach to how the utility function is used as a control scheme for the sensor to locate the target. Through simulations we show the proposed control policy is able to locate and track a moving target accurately and efficiently.**

I. Introduction

Tracking a moving object is an active research problem which has various applications in different industries such as astronautical [1] and biomedical [2]. In these applications, information must be gathered quickly and efficiently. One way to go about this is to gather information about the environment and target, and determine the best path for the sensor based on the information gathered. This paper focuses on efficiently estimating the location of a moving target in an unknown environment through active exploration, without the computational cost of mapping, with limited information provided by simple sensors (e.g. range-only, bearing-only, etc.). This problem can be thought of as an active perception problem in which one must find a control policy that is optimal in an information-theoretic sense [3]. Minimizing entropy, or equivalently, maximizing information has been successfully used to gather measurements [4], [5]. Many information-based control laws are computationally tractable because of the greedy maximizations they take over one time step or are used offline. The problem becomes computationally expensive when we want to compute the information over a finite time horizon and between multiple sensors, due to the integration over the joint measurement space of the sensor.

We propose to develop techniques to allow sensors to make autonomous decisions on which actions they should take to successively explore an environment to track a moving target. These autonomous actions can be thought of as decisions made to acquire the most information and used to decide if the information acquired is sufficient enough to pursue the target. The fundamental challenge in designing these autonomous decision processes is that agents have to operate in the face of uncertainty about the state of the environment. Our intended functionality is to endow sensors with the ability to reason in the face of this uncertainty about what are actions they should take given the partial knowledge they have about the environment.

Most research done in information-based control policies uses an approximation of Shannon or Fisher information [3], [6], [7]. Our primary contribution to this field of control policies is a method based on using a Gaussian mixture representation with Rényi's quadratic entropy, a generalization of Shannon's entropy [8]. The combination of using Gaussian mixture representation and Rényi's quadratic entropy to quantify the uncertainty has a closed-form expression when updating the belief of the state, reducing the computation complexity. Extending the expression over a finite time horizon can be done by approximating the conditional quadratic entropy over the time horizon. By estimating the target's location using a Gaussian mixture model, we are able to quickly determine paths which will yield the most useful measurements in localizing the target's state. There are many advantages in using Gaussian mixtures, the first and most useful is that an extension to multiple targets becomes trivial: as long as the number of modes in the Gaussian mixture is greater than the number of targets, our approach will be able to efficiently track the targets' states. Since there is a closed-form expression for the conditional quadratic entropy for one time step, extending the expression to account for a finite time horizon can be done with Monte Carlo integration.

There is a plethora of work on information gathering. In [3] and [9], the authors developed controllers to locate static targets by looking at the rate of change of the Fisher information matrix. Their work assumes a Gaussian belief. The authors of [7] propose an algorithm to localize an unknown number of targets in an environment by approximating the mutual information between the sensors' measurements and targets' states estimated using finite set statistics. They focus on a one step time horizon.

The authors in [10] addressed the problem of controlling mobile sensors to improve gathering information autonomously by presenting a non-greedy centralized solution

and an equivalent decentralized control scheme for active SLAM. The authors in [11] developed a target tracking algorithm based on Rényi's entropy using a combination of a particle filter and Gaussian mixture model to represent the distributions of the target's state. This paper relies heavily on a particle filter approach to estimate the target's location, which can be computationally expensive. It also only focuses on a one step horizon.

Our approach is similar to [12], in which the authors developed a control policy for a small team of robots to estimate the location of multiple mobile targets using an approximation of Shannon's mutual information. One downside to this approach is scalability. The mutual information approximations made by these authors is only computationally tractable for small groups of sensors. Our approach simplifies the order of computation by using the closed-form expression of the conditional entropy resulting from using a Gaussian mixture representation.

In this work we focus on using a range-bearing sensor to localize and track a moving target. In Sec. III, we detail the problem formulation. The utility function used in the control policy is defined in Sec. IV. In Sec. V, we outline how the sensor uses the control policy for target tracking. Finally, we present simulations in Sec. VI to validate the control policy.

II. Related Work

There is a plethora of work on information gathering. In [3] and [9], the authors developed controllers to locate static targets by looking at the rate of change of the Fisher information matrix. Their work assumes a Gaussian belief. The authors of [7] propose an algorithm to localize an unknown number of targets in an environment by approximating the mutual information between the sensors' measurements and targets' states estimated using finite set statistics. They focus on a one step time horizon.

The authors in [10] addressed the problem of controlling mobile sensors to improve gathering information autonomously by presenting a non-greedy centralized solution and an equivalent decentralized control scheme for active SLAM. The authors in [11] developed a target tracking algorithm based on Rényi's entropy using a combination of a particle filter and Gaussian mixture model to represent the distributions of the target's state. This paper relies heavily on a particle filter approach to estimate the target's location, which can be computationally expensive. It also only focuses on a one step horizon.

Our approach is similar to [12], in which the authors developed a control policy for a small team of robots to estimate the location of multiple mobile targets using an approximation of Shannon's mutual information. One downside to this approach is scalability. The mutual information approximations made by these authors is only computationally tractable for small groups of sensors. Our approach simplifies the order of computation by using the closed-form expression of the conditional entropy resulting from using a Gaussian mixture representation.

III. Problem Formulation

In this section we describe the problem of target tracking using an entropy based controller in more detail. The proposed problem involves one mobile sensor localizing and then tracking a moving target using only a simple sensor. Its control law is governed by minimizing the uncertainty of the target's state given the current measurement. For ease of reference, Table I, summarizes all important variables mentioned in this paper, while Figure 1 depicts the problem broken down into a block diagram.

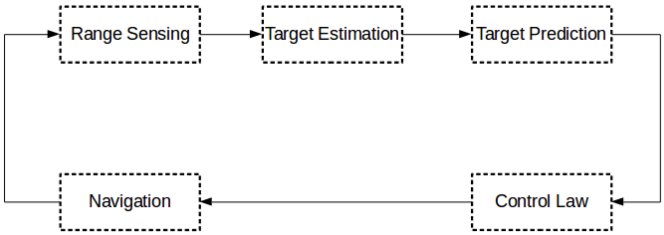


Fig. 1: Sensor Block Diagram

Consider a mobile sensor governed by the the following motion model:

$$x_{t+1} = f(x_t, u_t) \tag{1}$$

where $x_t \in \mathcal{X} \cong \mathbb{R}^{n_x}$ is the sensor state at time t, $u_t \in \mathcal{U}_t$ is the control input selected at time t, and $\mathcal{U}_t$ is a finite set of valid controls. The sensor must track the state of a target whose dynamics are governed by the linear target motion model:

$$y_{t+1} = Ay_t + \epsilon_t, \quad \varepsilon_t \sim \mathcal{N}(\varepsilon_t; 0, R) \tag{2}$$

where $y_t \in \mathbb{R}^{n_y}$ is the target state, $A \in \mathbb{R}^{n_y \times n_y}$, and ε_t is additive Gaussian noise. The sensor observes the state of the target through the nonlinear sensor observation model:

$$z_t = h(x_t, y_t) + \delta_t, \quad \delta_t \sim \mathcal{N}(\delta_t, 0, Q) \tag{3}$$

where $z_t \in \mathbb{R}^{n_z}$ is the measurement at time t, $h(x_t, y_t) \in \mathbb{R}^{n_z \times n_z}$ and δ_t is additive Gaussian noise. We assume the measurement noise is independent for any pairs of time as well as independent from the target motion model noise, ε_t.

Problem (Motion Planning Based on Uncertainty)**.** *Given the initial sensor position x_0, a prior distribution of the target state y_0, and a planning time horizon, T, the sensor is tasked to choose a sequence of control inputs $\hat{u}_{1:T-1} = (\hat{u}_1, \ldots, \hat{u}_{T-1})$, where $\hat{u}_i \in \mathcal{U}_i$, which minimizes the uncertainty of the final target state y_T.*

$$\begin{aligned} \underset{\hat{u}_{1:T-1}}{\arg\min} \quad & \mathbb{H}(y_T \mid z_{1:T}) \\ s.t. \quad & x_{t+1} = f(x_t, \hat{u}_t), && t = 0, \ldots, T-1, \\ & y_{t+1} = Ay_t + \varepsilon_t, && t = 1, \ldots, T, \\ & z_t = h(x_t, y_t) + \delta_t, && t = 1, \ldots, T \end{aligned} \tag{4}$$

where $\mathbb{H}(\cdot)$ is a measure of the uncertainty.

TABLE I: Notation

Symbol	Description
t	Current time
x	Sensor state
y	Target state
z	Measurement
u	Control inputs
T	Time horizon
M	Number of Gaussian modes in mixture model
N	Number of samples in Monte Carlo integration

IV. Rényi's Quadratic Entropy

In this section we will define the measure of uncertainty used to design the sensor's controller. There are different metrics for quantifying the uncertainty of a target state. Shannon's entropy is a widely used measure for uncertainty [13], and was generalized by Rényi in [8] as a parameter of α. Rényi's quadratic entropy ($\alpha = 2$) is defined by

$$\mathbb{H}_2(X) = -\log \int p(x)^2 dx \tag{5}$$

This expression turns out to be very useful when the distribution, $p(x)$, is a Gaussian mixture. Specifically, authors in [14] found that the quadratic entropy for a Gaussian mixture random variable, X, with weights, means, and covariance matrices, w_i, μ_i, and Σ_i respectively for $i = \{1, 2, \ldots, M\}$ has the following closed-form expression:

$$\mathbb{H}_2(X) = -\log \left[\sum_{i=1}^{M} \sum_{j=1}^{M} w_i w_j \cdot \phi(\mu_i; \mu_j, \Sigma_i + \Sigma_j) \right] \tag{6}$$

In this paper, we exploit this closed-form expression by using Rényi's quadratic entropy to quantify the uncertainty of the final target state when given a set of measurements,

$$\begin{aligned} \mathbb{H}_2(Y_T \mid Z_{1:T}) &= \int p(z_{1:T}) \mathbb{H}_2(y_T \mid Z_{1:T} = z_{1:T}) dz_{1:T} \\ &= -\int p(z_{1:T}) \log \int p(y_T \mid z_{1:T})^2 dy_T dz_{1:T} \end{aligned} \tag{7}$$

V. General Approach

In order to quantify the entropy of the target state, we must first find an expression for the state of the target, $p(y_T | z_{1:T})$. If the prior distribution of the target is assumed to be a Gaussian mixture model defined by

$$f_0(y_0 \mid w_0, \mu_0, \Sigma_0) = \sum_{i=1}^{M} w_i^0 \cdot \phi(y_0; \mu_i^0, \Sigma_i^0) \tag{8}$$

where M denotes the number of modes in the Gaussian mixture, $w_0 = [w_1^0, \ldots, w_M^0]$, $\mu_0 = [\mu_1^0, \ldots, \mu_M^0]$, and $\Sigma_0 = [\Sigma_1^0, \ldots, \Sigma_M^0]$ are the weights, mean, and covariance matrices of the individual Gaussian modes, respectively, then the belief of the target is also a Gaussian mixture.

Since the prior and belief are both defined by Gaussian mixture models, the belief can be updated using a Gaussian mixture filter outlined by [15]. The Gaussian mixture filter is based on a typical recursive Bayesian filter [16], where the first step uses the process model to predict the target's state and the second step uses the measurement to update the belief of the state.

The prediction step of the Gaussian mixture filter updates the belief of the state based on the dynamics of the target's motion in (2). The details of the prediction step are deferred to the Appendix.

Predict:

$$\begin{aligned} \overline{bel}(y_t) &= p(y_t \mid z_{1:t-1}) \\ &= \int p(y_t \mid y_{t-1}) bel(y_{t-1}) dy_{t-1} \\ &= \sum_{i=1}^{M} w_{t,i}^p \phi(y_t; \mu_{t,i}^p, \Sigma_{t,i}^p) \end{aligned} \tag{9}$$

where by $w_{t,i}^p = w_{t-1,i}^e$, $\mu_{t,i}^p = A\mu_{t-1,i}^e$, and $\Sigma_{t,i}^p = A\Sigma_{t-1,i}^e A^\top + R$.

The measurement update step of the Gaussian mixture filter updates the belief of the state based on the measurement model given in (3). The details of the measurement step are deferred to the Appendix.

Update:

$$\begin{aligned} bel(y_t) &= p(y_t \mid z_{1:t}) \\ &= c_t p(z_t \mid y_t) \overline{bel}(y_t) \\ &= \sum_{i=1}^{M} w_{t,i}^e \phi(y_t; \mu_{t,i}^e, \Sigma_{t,i}^e) \end{aligned} \tag{10}$$

where the new weights, means, and covariance matrices are detailed in (11). In (11), $H(x_t, \mu_{t,i}^e)$ is the Jacobian of $h(x_t, y_t)$ evaluated at $y_t = \mu_{t,i}^e$.

$$w_{t,i}^e \triangleq c_t w_{t,i}^p \phi(z_t; \mu_{t,i}^z, \Sigma_{t,i}^z) \tag{11a}$$

$$\mu_{t,i}^e = h(x_t, \mu_{t,i}^p) \tag{11b}$$

$$\Sigma_{t,i}^e = (\mathbb{I} - K_t H(x_t, \mu_{t,i}^p)) \Sigma_{t,i}^p \tag{11c}$$

$$K_t = \Sigma_{t,i}^p H(x_t, \mu_{t,i}^e)^\top \left[H(x_t, \mu_{t,i}^e) \Sigma_{t,i}^p H(x_t, \mu_{t,i}^e)^\top + Q \right]^{-1} \tag{11d}$$

Notice, since the prior and updated belief of the state is always a Gaussian mixture, we can use the result from [14], summarized in (6) to reduce the inner integral in (7).

$$\begin{aligned} &\mathbb{H}_2(Y_T \mid Z_{1:T}) \\ &= -\int p(z_{1:T}) \log \left[\sum_{i,j=1}^{M} w_{T,i}^e w_{T,j}^e \phi(\mu_{T,i}^e; \mu_{T,j}^e, \Sigma_{T,i}^e + \Sigma_{T,j}^e) \right] \end{aligned} \tag{12}$$

In order to find an expression for the conditional entropy, we need to find an expression for the measurement distribution, $p(z_{1:T})$. This can be done by expanding the expression using Bayes' rule and solving the integral using the recursive Bayesian filter equations. The details are deferred to the Appendix.

$$
\begin{aligned}
p(z_{1:T}) &= \int p(z_{1:T}, y_{0:T}) dy_{0:T} \\
&= \int p(z_{1:T} \mid p(y_{0:T})) p(y_{0:T}) dy_{0:T} \\
&= \int \prod_{t=1}^{T} \phi(z_t; h(x_t, y_t), Q) \prod_{t=1}^{T} p(y_t \mid y_{t-1}) p(y_0) dy_{0:T} \\
&= \sum_{i=1}^{M} w_i^z \phi\left(z_{1:T}; \mu_i^z, \Sigma_i^z\right)
\end{aligned} \tag{13}
$$

where the new weights, means, and covariance matrices are detailed in (14).

$$w_i^z = w_i^0 \tag{14a}$$

$$\mu_i^z = h(x_{1:T}, A_{1:T}\mu_i^0) \tag{14b}$$

$$\Sigma_i^z = H(x_T, \mu_i^0)\left[A_{1:T}\Sigma_i^0 A_{1:T}^\top + R_{1:T}\right] H(x_T, \mu_i^0)^\top + Q_{1:T} \tag{14c}$$

To resolve the last integral in (7), we can use Monte Carlo integration, a technique to numerically compute integrals using random variables. This approximation will increase the complexity of the calculation based on the number of samples used in the integral approximation, but since the measurement distribution has a closed-form expression, sampling can be done efficiently.

Monte Carlo Integration:

Let $p(z_{1:T}) \sim \sum_{i=1}^{M} w_i^z \phi\left(z_{1:T}; \mu_i^z, \Sigma_i^z\right)$

1) Repeat N times:
   a) Sample $k \sim [w_1^z, \cdots, w_M^z]$
   b) Sample $z_{1:T}^{(j)} \sim \phi\left(z_{1:T}; \mu_k^z, \Sigma_k^z\right)$
2) $H_2(Y_T \mid Z_{1:T}) \approx \frac{1}{N} \sum_{j=1}^{N} H_2\left(Y_T \mid z_{1:T}^{(j)}\right)$

Now we can approximate (7), with a closed for expression.

$$
\begin{aligned}
H_2(Y_T \mid Z_{1:T}) &\approx \frac{1}{N} \sum_{j=1}^{N} H_2\left(Y_T \mid z_{1:T}^{(j)}\right) \\
&= -\frac{1}{N} \sum_{j=1}^{N} \log \left[\sum_{i,j=1}^{M} w_{T,i}^e w_{T,j}^e \cdot \right. \\
&\qquad \left. \phi\left(\mu_{T,i}^e; \mu_{T,j}^e, \Sigma_{T,i}^e + \Sigma_{T,j}^e\right) \right]
\end{aligned} \tag{15}
$$

VI. Results

In this section we will present two simulations to evaluate the proposed quadratic entropy based controller. These simulations were done in MATLAB. The first simulation compares the uncertainty metric presented in Section V to the information metric presented in [12].

Figure 2 was generated by allowing the Gaussian mixture prior of the system to start with means distributed in a circle centered at the origin depicted by the black diamonds. The modes have equal weights. The quadratic entropy was calculated at grid points, then interpolated to get the contour plots across the entire environment. As illustrated in Figure 2 the maximum entropy, corresponding to the most uncertainty, or equivalently the least information, is in the center of the means. The results presented in this figure directly correspond to the results presented in Figure 5(f) of [12], in which the approximate Shannon's mutual information of the target's belief is lowest in the center of the 10 particles distributed in a circle.

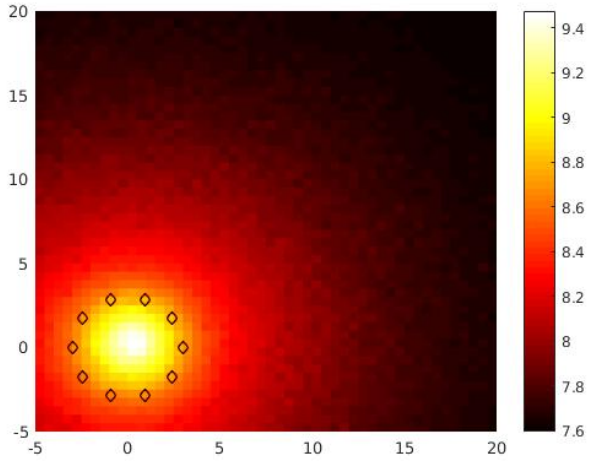


Fig. 2: Approximating Rényi's quadratic entropy with a circular distributed Gaussian mixture prior

To further evaluate the proposed uncertainty measure, we ran a simulation to track a moving target following the process outlined in Section V. A range-only sensor was used as the sensing modality, i.e. (3) can be rewritten as

$$z_t = \|x_t - y_t\|_2 + \delta_t, \quad \delta_t \sim \mathcal{N}(\delta_t, 0, Q) \tag{16}$$

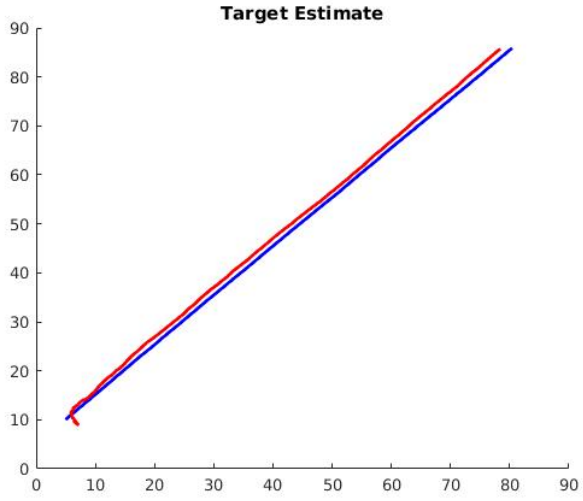


Fig. 3: Estimating target's location using Rényi's quadratic entropy based controller

Figure 3 shows one run of the simulation. There is some discrepancy between the actual target's location represented

by the blue line and the estimated target's location represented by the red line. This is due to the noise in the measurement. When Q is reduced, the discrepancy between the actual and estimated target's location decreases as expected.

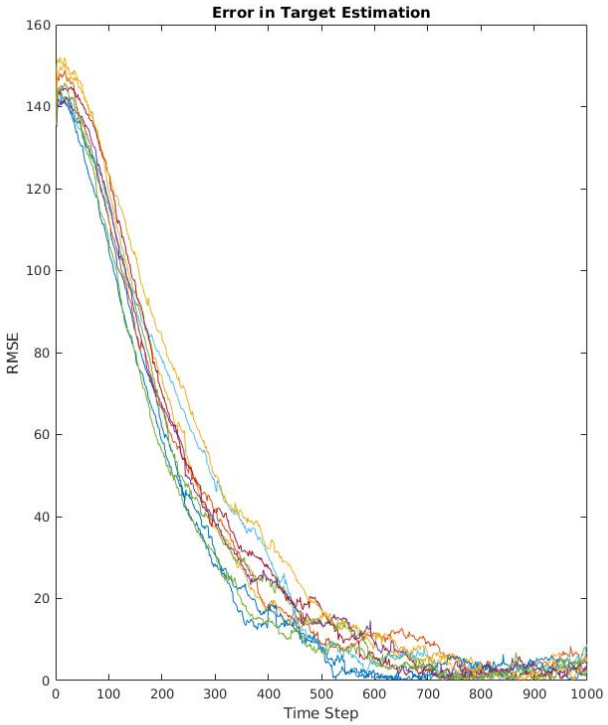


Fig. 4: RMSE of target estimator using Rényi's quadratic entropy based controller for ten simulations.

Figure 4 depicts the root mean squared error of ten simulation runs of the experiment described above. As seen, the estimator is able to quickly locate the target's position and track it with a relatively low RMSE. The RMSE depends directly on the measurement error, Q, as well as how well the target follows the predicted process model.

VII. Conclusion

This research complements the work of Dr. E. Stump and Dr. B. Sadler at ARL on developing control schemes based on mutual information to localize targets without continuous coordination of agents. Continuous coordination consumes resources and limits the mobility of the team, while decentralized control may produce redundancy in exploration. A combination of both reduces the overhead of coordination while allowing agents to explore a broader area.

In this paper we presented a controller which minimizes the Rényi's quadratic entropy over a finite time horizon allowing a mobile sensor to track a moving target in an unknown environment. We exploited the closed-form expression of Rényi's quadratic entropy for a Gaussian mixture distribution. We then used an approximation to integrate the utility function over multiple time steps to formulate an expression for the uncertainty over a finite time horizon. We validated our approach by comparing it to results presented in [12] and running a simulation in which the sensor successfully tracked a moving target. There is still a lot of work needed to be done to complete this work in active exploration for target tracking. An indoor real world experiment still needs to be performed. In the experiment, we plan to track a moving person using a Scarab with a projective camera as the sensor. The software will be written in C++ and interfaced using the Robot Operating System (ROS).

There are many interesting directions for future research. We plan to also look at using the Cauchy-Schwartz definition of mutual information as an information measure. The Cauchy-Schwartz definition of mutual information also has a closed-form expression for Gaussian mixture models. We would like to compare the computational complexity of Rényi's quadratic entropy with Cauchy-Schwartz mutual information. We believe that the Monte Carlo integration done Section V may not be needed in the latter information measure. Since the latter measures information rather than uncertainty, it will be inversely proportional to the metric used in this paper.

We would also like to extend this to multiple targets and multiple sensors. As stated previously, extending this work to multiple targets should be trivial, as long as the number of modes, M, is larger than the number of targets, this algorithm should be able to track all the targets, while still being computationally tractable. This extension to multiple sensors allows us to formulate the problem in terms of a game. The sensors, or agents, will have to coordinate to track the targets. They must be able to decide whether communicating with other agents will provide more information than continuing to explore on their own. This requires another controller which decides between coordinating and exploring. An additional mode can be added to the new controller if a server-like entity is introduced into the game. The server acts as a history logger and stores information which has been uploaded by agents. Agents can also download servers, but must be physically near a server location to do this.

Another extension to this problem would be to investigate continuous trajectory planning for the sensors. Currently, the inputs, u_t, are points sensors can navigate to in one time step. We want to look at ways of interpolating the points to form continuous trajectories that also minimize the target's uncertainty.

Appendix

A. Prediction Step

$$\begin{aligned}\overline{bel}(y_t) &= \int p(y_t \mid y_{t-1}) bel(y_{t-1}) dy_{t-1} \\ &= \int \phi(y_t; Ay_{t-1}, R) \sum_{i=1}^{M} w_{t-1,i}^{e} \\ &\qquad \phi(y_{t-1}; \mu_{t-1,i}^{e}, \Sigma_{t-1,i}^{e}) dy_{t-1} \\ &= \sum_{i=1}^{M} w_{t,i}^{p} \phi(y_t; \mu_{t,i}^{p}, \Sigma_{t,i}^{p})\end{aligned}$$

where the weights, means, and covariance matrices are defined as

$$\begin{aligned} w^p_{t,i} &= w^e_{t-1,i} \\ \mu^p_{t,i} &= A\mu^e_{t-1,i} \\ \Sigma^p_{t,i} &= A\Sigma^e_{t-1,i}A^\top + R \end{aligned}$$

B. Measurement Step

$$\begin{aligned} bel(y_t) &= c_t p(z_t \mid y_t)\overline{bel}(y_t) \\ &= c_t\phi(z_t; h(x_t, y_t), Q)\sum_{i=1}^{M} w^p_{t,i}\phi(y_t; \mu^p_{t,i}, \Sigma^p_{t,i}) \\ &= \sum_{i=1}^{M} c_t w^p_{t,i}\phi(z_t; h(x_t, y_t), Q)\phi(y_t; \mu^p_{t,i}, \Sigma^p_{t,i}) \\ &= \sum_{i=1}^{M} c_t w^p_{t,i}\phi(z_t; \mu^z_{t,i}, \Sigma^z_{t,i})\cdot \\ &\qquad \frac{\phi(z_t; h(x_t, y_t), Q)\phi(y_t; \mu^p_{t,i}, \Sigma^p_{t,i})}{\phi(z_t; \mu^z_{t,i}, \Sigma^z_{t,i})} \\ &= \sum_{i=1}^{M} w^e_{t,i}\phi(y_t; \mu^e_{t,i}, \Sigma^e_{t,i}) \end{aligned}$$

where the weights, means, and covariance matrices are defined as

$$\begin{aligned} w^e_{t,i} &\triangleq c_t w^p_{t,i}\phi(z_t; \mu^z_{t,i}, \Sigma^z_{t,i}) \\ \mu^e_{t,i} &= h(x_t, \mu^p_{t,i}) \\ \Sigma^e_{t,i} &= (\mathbb{I} - K_t H(x_t, \mu^p_{t,i}))\Sigma^p_{t,i} \\ K_t &= \Sigma^p_{t,i}H(x_t, \mu^e_{t,i})^\top \left[H(x_t, \mu^e_{t,i})\Sigma^p_{t,i}H(x_t, \mu^e_{t,i})^\top + Q\right]^{-1} \end{aligned}$$

and $H(x_t, \mu^e_{t,i})$ is the Jacobian of $h(x_t, y_t)$ evaluated at $y_t = \mu^e_{t,i}$.

C. Measurement Distribution

$$\begin{aligned} p(z_t) = \frac{1}{c_t} &= \int p(z_t \mid y_t)\overline{bel}(y_t)dy_t \\ &= \int \phi(z_t; h(x_t, y_t), Q)\sum_{i=1}^{M} w^p_{t,i}\phi(y_t; \mu^p_{t,i}, \Sigma^p_{t,i})dy_t \\ &= \sum_{i=1}^{M} w^p_{t,i}\int \phi(z_t; h(x_t, y_t), Q)\phi(y_t; \mu^p_{t,i}, \Sigma^p_{t,i})dy_t \\ &= \sum_{i=1}^{M} w^p_{t,i}\phi(z_t; \mu^z_{t,i}, \Sigma^z_{t,i}) \end{aligned}$$

where

$$\begin{aligned} \mu^z_{t,i} &= h(x_t, \mu^p_{t,i}) \\ \Sigma^z_{t,i} &= H(x_t, \mu^p_{t,i})\Sigma^p_{t,i}H(x_t, \mu^p_{t,i})^\top + Q \end{aligned}$$

D. Measurement Distribution Over a Time Horizon

$$\begin{aligned} p(z_{1:T}) &= \int p(z_{1:T}, y_{0:T})dy_{0:T} \\ &= \int p(z_{1:T} \mid p(y_{0:T}))p(y_{0:T})dy_{0:T} \\ &= \int \prod_{t=1}^{T}\phi(z_t; h(x_t, y_t), Q)\prod_{t=1}^{T} p(y_t \mid y_{t-1}) \\ &\qquad p(y_0)dy_{0:T} \\ &= \int \prod_{t=1}^{T}\phi(z_t; h(x_t, y_t), Q)\prod_{t=1}^{T}\phi(y_t; Ay_{t-1}, R) \\ &\qquad \sum_{i=1}^{M} w^0_i\phi(y_0; \mu^0_i, \Sigma^0_i)dy_{0:T} \\ &= \sum_{i=1}^{M} w^0_i\int \phi(z_{1:T}; h(x_{1:T}, y_{1:T}), \tilde{Q})\phi(y_{1:T}; \tilde{A}y_0, \tilde{R}) \\ &\qquad \sum_{i=1}^{M} w^0_i\phi(y_0; \mu^0_i, \Sigma^0_i)dy_{0:T} \\ &= \sum_{i=1}^{M} w^z_i\phi\left(z_{1:T}; \mu^z_i, \Sigma^z_i\right) \end{aligned}$$

where

$$\begin{aligned} w^z_i &= w^0_i \\ \mu^z_i &= h(x_{1:T}, A_{1:T}\mu^0_i) \\ \Sigma^z_i &= H(x_T, \mu^0_i)\left[A_{1:T}\Sigma^0_i A^\top_{1:T} + R_{1:T}\right]H(x_T, \mu^0_i)^\top + Q_{1:T} \end{aligned}$$

VIII. Technology Transitions

The techniques developed through this thrust correspond the MAST programs mission to develop the technology to enable true autonomous operation of each member of the team.

References

[1] D. Mehrholz, L. Leushacke, W. Flury, R. Jehn, H. Klinkrad, and M. Landgraf, "Detecting, tracking and imaging space debris," *ESA Bulletin(0376-4265)*, no. 109, pp. 128–134, 2002.

[2] S. T. Acton and N. Ray, "Biomedical image analysis: tracking," *Synthesis Lectures on Image, Video, and Multimedia Processing*, vol. 2, no. 1, pp. 1–152, 2006.

[3] B. Grocholsky, "Information-theoretic control of multiple sensor platforms," Ph.D. dissertation, Citeseer, 2002.

[4] G. M. Hoffmann and C. J. Tomlin, "Mobile sensor network control using mutual information methods and particle filters," *IEEE Transactions on Automatic Control*, vol. 55, no. 1, pp. 32–47, 2010.

[5] B. J. Julian, M. Angermann, M. Schwager, and D. Rus, "A scalable information theoretic approach to distributed robot coordination," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 5187–5194.

[6] B. Charrow, "Information-Theoretic Active Perception for Multi-Robot Teams," Ph.D. dissertation, University of Pennsylvania, 2015.

[7] M. Schwager, P. Dames, D. Rus, and V. Kumar, "A multi-robot control policy for information gathering in the presence of unknown hazards," in *International Symposium on Robotics Research*, 2011, pp. 17–33.

[8] A. Rényi, "On measures of entropy and information," in *Fourth Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, 1961, pp. 547–561.

[9] E. Stump, V. Kumar, B. Grocholsky, and P. M. Shiroma, "Control for localization of targets using range-only sensors," *The International Journal of Robotics Research*, vol. 28, no. 6, pp. 743–757, 2009.
[10] N. Atanasov, J. Le Ny, K. Daniilidis, and G. Pappas, "Decentralized Active Information Acquisition: Theory and Application to Multi-Robot SLAM," in *IEEE Int. Conf. on Robotics and Automation*, 2015, pp. 4775–4782. [Online]. Available: http://dx.doi.org/10.1109/ICRA.2015.7139863
[11] D. Gu and H. Hu, "Rényi entropy based target tracking in mobile sensor networks," *IFAC Proceedings Volumes*, vol. 44, no. 1, pp. 13 558–13 563, 2011.
[12] B. Charrow, V. Kumar, and N. Michael, "Approximate representations for multi-robot control policies that maximize mutual information," *Autonomous Robots*, vol. 37, no. 4, pp. 383–400, 2014.
[13] T. M. Cover and J. A. Thomas, "Elements of information theory 2nd edition," 2006.
[14] F. Wang, T. Syeda-Mahmood, B. C. Vemuri, D. Beymer, and A. Rangarajan, "Closed-form jensen-renyi divergence for mixture of gaussians and applications to group-wise shape registration," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2009, pp. 648–655.
[15] M. Huber, *Nonlinear Gaussian Filtering: Theory, Algorithms, and Applications*. KIT Scientific Publishing, 2015, vol. 19.
[16] S. Thrun, W. Burgard, and D. Fox, "Probabilistic robotics," 2005.

Sensor-based Motion Planning using Gaussian Mixtures and Rényi's Quadratic Entropy

CNC 16-17-2.1 Distributed Information Games for Autonomous Team Behavior

Shih-Ling Phuong, Nikolay Atanasov, Alejandro Ribeiro
Dept. of Electrical and Systems Engineering,
University of Pennsylvania,
Philadelphia, Pennsylvania,
{sphuong, atanasov, aribeiro}@seas.upenn.edu.

***Abstract*—This paper examines the problem of tracking a moving target by exploring an unknown environment using a mobile sensor. The sensor actively explores without mapping by modeling the state of the target as a Gaussian mixture model. This paper focuses on designing a control policy for the sensor which minimizes a utility function based on the uncertainty of the target's state. First, we present an information-theoretic utility function derived using Rényi's definition of entropy to quantify the uncertainty of the target's state given the sensor measurements. We then outline the genearl approach to how the utility function is used as a control scheme for the sensor to locate the target. Through simulations we show the proposed control policy is able to locate and track a moving target accurately and efficiently.**

I. Introduction

Tracking a moving object is an active research problem which has various applications in different industries such as astronautical [1] and biomedical [2]. In these applications, information must be gathered quickly and efficiently. One way to go about this is to gather information about the environment and target, and determine the best path for the sensor based on the information gathered. This paper focuses on efficiently estimating the location of a moving target in an unknown environment through active exploration, without the computational cost of mapping, with limited information provided by simple sensors (e.g. range-only, bearing-only, etc.). This problem can be thought of as an active perception problem in which one must find a control policy that is optimal in an information-theoretic sense [3]. Minimizing entropy, or equivalently, maximizing information has been successfully used to gather measurements [4], [5]. Many information-based control laws are computationally tractable because of the greedy maximizations they take over one time step or are used offline. The problem becomes computationally expensive when we want to compute the information over a finite time horizon and between multiple sensors, due to the integration over the joint measurement space of the sensor.

We propose to develop techniques to allow sensors to make autonomous decisions on which actions they should take to successively explore an environment to track a moving target. These autonomous actions can be thought of as decisions made to acquire the most information and used to decide if the information acquired is sufficient enough to pursue the target. The fundamental challenge in designing these autonomous decision processes is that agents have to operate in the face of uncertainty about the state of the environment. Our intended functionality is to endow sensors with the ability to reason in the face of this uncertainty about what are actions they should take given the partial knowledge they have about the environment.

Most research done in information-based control policies uses an approximation of Shannon or Fisher information [3], [6], [7]. Our primary contribution to this field of control policies is a method based on using a Gaussian mixture representation with Rényi's quadratic entropy, a generalization of Shannon's entropy [8]. The combination of using Gaussian mixture representation and Rényi's quadratic entropy to quantify the uncertainty has a closed-form expression when updating the belief of the state, reducing the computation complexity. Extending the expression over a finite time horizon can be done by approximating the conditional quadratic entropy over the time horizon. By estimating the target's location using a Gaussian mixture model, we are able to quickly determine paths which will yield the most useful measurements in localizing the target's state. There are many advantages in using Gaussian mixtures, the first and most useful is that an extension to multiple targets becomes trivial: as long as the number of modes in the Gaussian mixture is greater than the number of targets, our approach will be able to efficiently track the targets' states. Since there is a closed-form expression for the conditional quadratic entropy for one time step, extending the expression to account for a finite time horizon can be done with Monte Carlo integration.

There is a plethora of work on information gathering. In [3] and [9], the authors developed controllers to locate static targets by looking at the rate of change of the Fisher information matrix. Their work assumes a Gaussian belief. The authors of [7] propose an algorithm to localize an unknown number of targets in an environment by approximating the mutual information between the sensors' measurements and targets' states estimated using finite set statistics. They focus on a one step time horizon.

The authors in [10] addressed the problem of controlling mobile sensors to improve gathering information autonomously by presenting a non-greedy centralized solution

and an equivalent decentralized control scheme for active SLAM. The authors in [11] developed a target tracking algorithm based on Rényi's entropy using a combination of a particle filter and Gaussian mixture model to represent the distributions of the target's state. This paper relies heavily on a particle filter approach to estimate the target's location, which can be computationally expensive. It also only focuses on a one step horizon.

Our approach is similar to [12], in which the authors developed a control policy for a small team of robots to estimate the location of multiple mobile targets using an approximation of Shannon's mutual information. One downside to this approach is scalability. The mutual information approximations made by these authors is only computationally tractable for small groups of sensors. Our approach simplifies the order of computation by using the closed-form expression of the conditional entropy resulting from using a Gaussian mixture representation.

In this work we focus on using a range-bearing sensor to localize and track a moving target. In Sec. III, we detail the problem formulation. The utility function used in the control policy is defined in Sec. IV. In Sec. V, we outline how the sensor uses the control policy for target tracking. Finally, we present simulations in Sec. VI to validate the control policy.

II. Related Work

There is a plethora of work on information gathering. In [3] and [9], the authors developed controllers to locate static targets by looking at the rate of change of the Fisher information matrix. Their work assumes a Gaussian belief. The authors of [7] propose an algorithm to localize an unknown number of targets in an environment by approximating the mutual information between the sensors' measurements and targets' states estimated using finite set statistics. They focus on a one step time horizon.

The authors in [10] addressed the problem of controlling mobile sensors to improve gathering information autonomously by presenting a non-greedy centralized solution and an equivalent decentralized control scheme for active SLAM. The authors in [11] developed a target tracking algorithm based on Rényi's entropy using a combination of a particle filter and Gaussian mixture model to represent the distributions of the target's state. This paper relies heavily on a particle filter approach to estimate the target's location, which can be computationally expensive. It also only focuses on a one step horizon.

Our approach is similar to [12], in which the authors developed a control policy for a small team of robots to estimate the location of multiple mobile targets using an approximation of Shannon's mutual information. One downside to this approach is scalability. The mutual information approximations made by these authors is only computationally tractable for small groups of sensors. Our approach simplifies the order of computation by using the closed-form expression of the conditional entropy resulting from using a Gaussian mixture representation.

III. Problem Formulation

In this section we describe the problem of target tracking using an entropy based controller in more detail. The proposed problem involves one mobile sensor localizing and then tracking a moving target using only a simple sensor. Its control law is governed by minimizing the uncertainty of the target's state given the current measurement. For ease of reference, Table I, summarizes all important variables mentioned in this paper, while Figure 1 depicts the problem broken down into a block diagram.

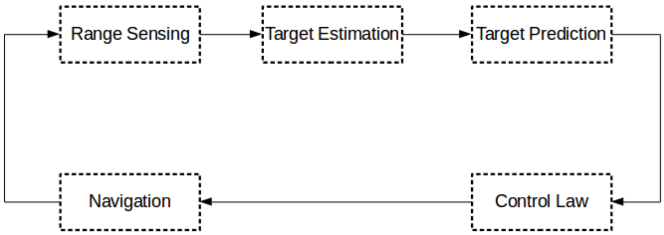


Fig. 1: Sensor Block Diagram

Consider a mobile sensor governed by the the following motion model:

$$x_{t+1} = f(x_t, u_t) \tag{1}$$

where $x_t \in \mathcal{X} \cong \mathbb{R}^{n_x}$ is the sensor state at time t, $u_t \in \mathcal{U}_t$ is the control input selected at time t, and $\mathcal{U}_t$ is a finite set of valid controls. The sensor must track the state of a target whose dynamics are governed by the linear target motion model:

$$y_{t+1} = Ay_t + \epsilon_t, \quad \varepsilon_t \sim \mathcal{N}(\varepsilon_t; 0, R) \tag{2}$$

where $y_t \in \mathbb{R}^{n_y}$ is the target state, $A \in \mathbb{R}^{n_y \times n_y}$, and ε_t is additive Gaussian noise. The sensor observes the state of the target through the nonlinear sensor observation model:

$$z_t = h(x_t, y_t) + \delta_t, \quad \delta_t \sim \mathcal{N}(\delta_t, 0, Q) \tag{3}$$

where $z_t \in \mathbb{R}^{n_z}$ is the measurement at time t, $h(x_t, y_t) \in \mathbb{R}^{n_z \times n_z}$ and δ_t is additive Gaussian noise. We assume the measurement noise is independent for any pairs of time as well as independent from the target motion model noise, ε_t.

Problem (Motion Planning Based on Uncertainty)**.** *Given the initial sensor position x_0, a prior distribution of the target state y_0, and a planning time horizon, T, the sensor is tasked to choose a sequence of control inputs $\hat{u}_{1:T-1} = (\hat{u}_1, \ldots, \hat{u}_{T-1})$, where $\hat{u}_i \in \mathcal{U}_i$, which minimizes the uncertainty of the final target state y_T.*

$$\begin{aligned} \underset{\hat{u}_{1:T-1}}{\arg\min} \quad & \mathbb{H}(y_T \mid z_{1:T}) \\ s.t. \quad & x_{t+1} = f(x_t, \hat{u}_t), && t = 0, \ldots, T-1, \\ & y_{t+1} = Ay_t + \varepsilon_t, && t = 1, \ldots, T, \\ & z_t = h(x_t, y_t) + \delta_t, && t = 1, \ldots, T \end{aligned} \tag{4}$$

where $\mathbb{H}(\cdot)$ is a measure of the uncertainty.

TABLE I: Notation

Symbol	Description
t	Current time
x	Sensor state
y	Target state
z	Measurement
u	Control inputs
T	Time horizon
M	Number of Gaussian modes in mixture model
N	Number of samples in Monte Carlo integration

IV. Rényi's Quadratic Entropy

In this section we will define the measure of uncertainty used to design the sensor's controller. There are different metrics for quantifying the uncertainty of a target state. Shannon's entropy is a widely used measure for uncertainty [13], and was generalized by Rényi in [8] as a parameter of α. Rényi's quadratic entropy ($\alpha = 2$) is defined by

$$\mathbb{H}_2(X) = -\log \int p(x)^2 dx \tag{5}$$

This expression turns out to be very useful when the distribution, $p(x)$, is a Gaussian mixture. Specifically, authors in [14] found that the quadratic entropy for a Gaussian mixture random variable, X, with weights, means, and covariance matrices, w_i, μ_i, and Σ_i respectively for $i = \{1, 2, \ldots, M\}$ has the following closed-form expression:

$$\mathbb{H}_2(X) = -\log \left[\sum_{i=1}^{M} \sum_{j=1}^{M} w_i w_j \cdot \phi(\mu_i; \mu_j, \Sigma_i + \Sigma_j) \right] \tag{6}$$

In this paper, we exploit this closed-form expression by using Rényi's quadratic entropy to quantify the uncertainty of the final target state when given a set of measurements,

$$\begin{aligned} \mathbb{H}_2(Y_T \mid Z_{1:T}) &= \int p(z_{1:T}) \mathbb{H}_2(y_T \mid Z_{1:T} = z_{1:T}) dz_{1:T} \\ &= -\int p(z_{1:T}) \log \int p(y_T \mid z_{1:T})^2 dy_T dz_{1:T} \end{aligned} \tag{7}$$

V. General Approach

In order to quantify the entropy of the target state, we must first find an expression for the state of the target, $p(y_T | z_{1:T})$. If the prior distribution of the target is assumed to be a Gaussian mixture model defined by

$$f_0(y_0 \mid w_0, \mu_0, \Sigma_0) = \sum_{i=1}^{M} w_i^0 \cdot \phi(y_0; \mu_i^0, \Sigma_i^0) \tag{8}$$

where M denotes the number of modes in the Gaussian mixture, $w_0 = [w_1^0, \ldots, w_M^0]$, $\mu_0 = [\mu_1^0, \ldots, \mu_M^0]$, and $\Sigma_0 = [\Sigma_1^0, \ldots, \Sigma_M^0]$ are the weights, mean, and covariance matrices of the individual Gaussian modes, respectively, then the belief of the target is also a Gaussian mixture.

Since the prior and belief are both defined by Gaussian mixture models, the belief can be updated using a Gaussian mixture filter outlined by [15]. The Gaussian mixture filter is based on a typical recursive Bayesian filter [16], where the first step uses the process model to predict the target's state and the second step uses the measurement to update the belief of the state.

The prediction step of the Gaussian mixture filter updates the belief of the state based on the dynamics of the target's motion in (2). The details of the prediction step are deferred to the Appendix.

Predict:

$$\begin{aligned} \overline{bel}(y_t) &= p(y_t \mid z_{1:t-1}) \\ &= \int p(y_t \mid y_{t-1}) bel(y_{t-1}) dy_{t-1} \\ &= \sum_{i=1}^{M} w_{t,i}^p \phi(y_t; \mu_{t,i}^p, \Sigma_{t,i}^p) \end{aligned} \tag{9}$$

where by $w_{t,i}^p = w_{t-1,i}^e$, $\mu_{t,i}^p = A\mu_{t-1,i}^e$, and $\Sigma_{t,i}^p = A\Sigma_{t-1,i}^e A^\top + R$.

The measurement update step of the Gaussian mixture filter updates the belief of the state based on the measurement model given in (3). The details of the measurement step are deferred to the Appendix.

Update:

$$\begin{aligned} bel(y_t) &= p(y_t \mid z_{1:t}) \\ &= c_t p(z_t \mid y_t) \overline{bel}(y_t) \\ &= \sum_{i=1}^{M} w_{t,i}^e \phi(y_t; \mu_{t,i}^e, \Sigma_{t,i}^e) \end{aligned} \tag{10}$$

where the new weights, means, and covariance matrices are detailed in (11). In (11), $H(x_t, \mu_{t,i}^e)$ is the Jacobian of $h(x_t, y_t)$ evaluated at $y_t = \mu_{t,i}^e$.

$$w_{t,i}^e \triangleq c_t w_{t,i}^p \phi(z_t; \mu_{t,i}^z, \Sigma_{t,i}^z) \tag{11a}$$

$$\mu_{t,i}^e = h(x_t, \mu_{t,i}^p) \tag{11b}$$

$$\Sigma_{t,i}^e = (\mathbb{I} - K_t H(x_t, \mu_{t,i}^p)) \Sigma_{t,i}^p \tag{11c}$$

$$K_t = \Sigma_{t,i}^p H(x_t, \mu_{t,i}^e)^\top \left[H(x_t, \mu_{t,i}^e) \Sigma_{t,i}^p H(x_t, \mu_{t,i}^e)^\top + Q \right]^{-1} \tag{11d}$$

Notice, since the prior and updated belief of the state is always a Gaussian mixture, we can use the result from [14], summarized in (6) to reduce the inner integral in (7).

$$\begin{aligned} &\mathbb{H}_2(Y_T \mid Z_{1:T}) \\ &= -\int p(z_{1:T}) \log \left[\sum_{i,j=1}^{M} w_{T,i}^e w_{T,j}^e \phi(\mu_{T,i}^e; \mu_{T,j}^e, \Sigma_{T,i}^e + \Sigma_{T,j}^e) \right] \end{aligned} \tag{12}$$

In order to find an expression for the conditional entropy, we need to find an expression for the measurement distribution, $p(z_{1:T})$. This can be done by expanding the expression using Bayes' rule and solving the integral using the recursive Bayesian filter equations. The details are deferred to the Appendix.

$$
\begin{aligned}
p(z_{1:T}) &= \int p(z_{1:T}, y_{0:T}) dy_{0:T} \\
&= \int p(z_{1:T} \mid p(y_{0:T})) p(y_{0:T}) dy_{0:T} \\
&= \int \prod_{t=1}^{T} \phi(z_t; h(x_t, y_t), Q) \prod_{t=1}^{T} p(y_t \mid y_{t-1}) p(y_0) dy_{0:T} \\
&= \sum_{i=1}^{M} w_i^z \phi\left(z_{1:T}; \mu_i^z, \Sigma_i^z\right)
\end{aligned} \tag{13}
$$

where the new weights, means, and covariance matrices are detailed in (14).

$$w_i^z = w_i^0 \tag{14a}$$

$$\mu_i^z = h(x_{1:T}, A_{1:T}\mu_i^0) \tag{14b}$$

$$\Sigma_i^z = H(x_T, \mu_i^0)\left[A_{1:T}\Sigma_i^0 A_{1:T}^\top + R_{1:T}\right] H(x_T, \mu_i^0)^\top + Q_{1:T} \tag{14c}$$

To resolve the last integral in (7), we can use Monte Carlo integration, a technique to numerically compute integrals using random variables. This approximation will increase the complexity of the calculation based on the number of samples used in the integral approximation, but since the measurement distribution has a closed-form expression, sampling can be done efficiently.

Monte Carlo Integration:

Let $p(z_{1:T}) \sim \sum_{i=1}^{M} w_i^z \phi\left(z_{1:T}; \mu_i^z, \Sigma_i^z\right)$

1) Repeat N times:
   a) Sample $k \sim [w_1^z, \cdots, w_M^z]$
   b) Sample $z_{1:T}^{(j)} \sim \phi\left(z_{1:T}; \mu_k^z, \Sigma_k^z\right)$
2) $H_2(Y_T \mid Z_{1:T}) \approx \frac{1}{N} \sum_{j=1}^{N} H_2\left(Y_T \mid z_{1:T}^{(j)}\right)$

Now we can approximate (7), with a closed for expression.

$$
\begin{aligned}
H_2(Y_T \mid Z_{1:T}) &\approx \frac{1}{N} \sum_{j=1}^{N} H_2\left(Y_T \mid z_{1:T}^{(j)}\right) \\
&= -\frac{1}{N} \sum_{j=1}^{N} \log \left[\sum_{i,j=1}^{M} w_{T,i}^e w_{T,j}^e \cdot \right. \\
&\qquad \left. \phi\left(\mu_{T,i}^e; \mu_{T,j}^e, \Sigma_{T,i}^e + \Sigma_{T,j}^e\right) \right]
\end{aligned} \tag{15}
$$

VI. Results

In this section we will present two simulations to evaluate the proposed quadratic entropy based controller. These simulations were done in MATLAB. The first simulation compares the uncertainty metric presented in Section V to the information metric presented in [12].

Figure 2 was generated by allowing the Gaussian mixture prior of the system to start with means distributed in a circle centered at the origin depicted by the black diamonds. The modes have equal weights. The quadratic entropy was calculated at grid points, then interpolated to get the contour plots across the entire environment. As illustrated in Figure 2 the maximum entropy, corresponding to the most uncertainty, or equivalently the least information, is in the center of the means. The results presented in this figure directly correspond to the results presented in Figure 5(f) of [12], in which the approximate Shannon's mutual information of the target's belief is lowest in the center of the 10 particles distributed in a circle.

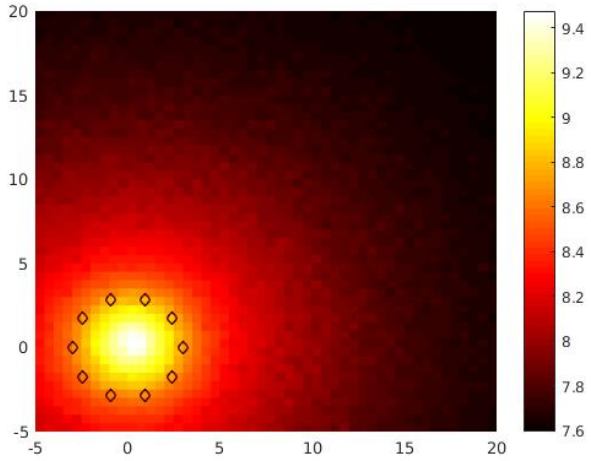


Fig. 2: Approximating Rényi's quadratic entropy with a circular distributed Gaussian mixture prior

To further evaluate the proposed uncertainty measure, we ran a simulation to track a moving target following the process outlined in Section V. A range-only sensor was used as the sensing modality, i.e. (3) can be rewritten as

$$z_t = \|x_t - y_t\|_2 + \delta_t, \quad \delta_t \sim \mathcal{N}(\delta_t, 0, Q) \tag{16}$$

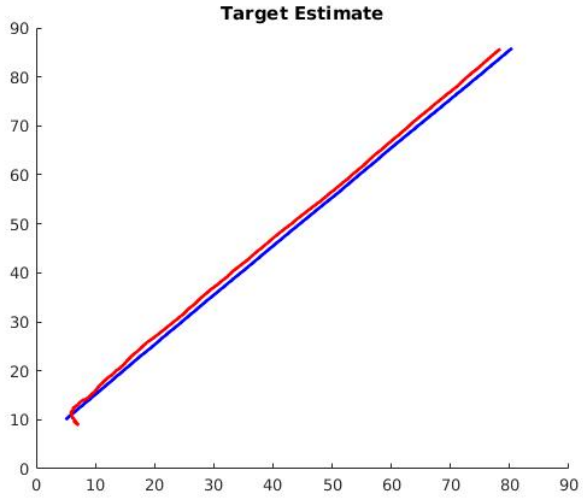


Fig. 3: Estimating target's location using Rényi's quadratic entropy based controller

Figure 3 shows one run of the simulation. There is some discrepancy between the actual target's location represented

by the blue line and the estimated target's location represented by the red line. This is due to the noise in the measurement. When Q is reduced, the discrepancy between the actual and estimated target's location decreases as expected.

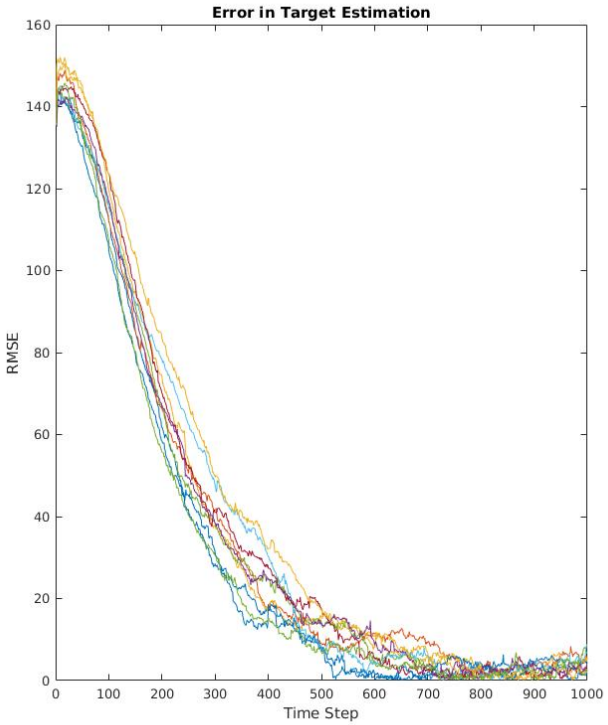


Fig. 4: RMSE of target estimator using Rényi's quadratic entropy based controller for ten simulations.

Figure 4 depicts the root mean squared error of ten simulation runs of the experiment described above. As seen, the estimator is able to quickly locate the target's position and track it with a relatively low RMSE. The RMSE depends directly on the measurement error, Q, as well as how well the target follows the predicted process model.

VII. Conclusion

This research complements the work of Dr. E. Stump and Dr. B. Sadler at ARL on developing control schemes based on mutual information to localize targets without continuous coordination of agents. Continuous coordination consumes resources and limits the mobility of the team, while decentralized control may produce redundancy in exploration. A combination of both reduces the overhead of coordination while allowing agents to explore a broader area.

In this paper we presented a controller which minimizes the Rényi's quadratic entropy over a finite time horizon allowing a mobile sensor to track a moving target in an unknown environment. We exploited the closed-form expression of Rényi's quadratic entropy for a Gaussian mixture distribution. We then used an approximation to integrate the utility function over multiple time steps to formulate an expression for the uncertainty over a finite time horizon. We validated our approach by comparing it to results presented in [12] and running a simulation in which the sensor successfully tracked a moving target. There is still a lot of work needed to be done to complete this work in active exploration for target tracking. An indoor real world experiment still needs to be performed. In the experiment, we plan to track a moving person using a Scarab with a projective camera as the sensor. The software will be written in C++ and interfaced using the Robot Operating System (ROS).

There are many interesting directions for future research. We plan to also look at using the Cauchy-Schwartz definition of mutual information as an information measure. The Cauchy-Schwartz definition of mutual information also has a closed-form expression for Gaussian mixture models. We would like to compare the computational complexity of Rényi's quadratic entropy with Cauchy-Schwartz mutual information. We believe that the Monte Carlo integration done Section V may not be needed in the latter information measure. Since the latter measures information rather than uncertainty, it will be inversely proportional to the metric used in this paper.

We would also like to extend this to multiple targets and multiple sensors. As stated previously, extending this work to multiple targets should be trivial, as long as the number of modes, M, is larger than the number of targets, this algorithm should be able to track all the targets, while still being computationally tractable. This extension to multiple sensors allows us to formulate the problem in terms of a game. The sensors, or agents, will have to coordinate to track the targets. They must be able to decide whether communicating with other agents will provide more information than continuing to explore on their own. This requires another controller which decides between coordinating and exploring. An additional mode can be added to the new controller if a server-like entity is introduced into the game. The server acts as a history logger and stores information which has been uploaded by agents. Agents can also download servers, but must be physically near a server location to do this.

Another extension to this problem would be to investigate continuous trajectory planning for the sensors. Currently, the inputs, u_t, are points sensors can navigate to in one time step. We want to look at ways of interpolating the points to form continuous trajectories that also minimize the target's uncertainty.

Appendix

A. Prediction Step

$$\begin{aligned}\overline{bel}(y_t) &= \int p(y_t \mid y_{t-1}) bel(y_{t-1}) dy_{t-1} \\ &= \int \phi(y_t; Ay_{t-1}, R) \sum_{i=1}^{M} w_{t-1,i}^{e} \\ &\qquad \phi(y_{t-1}; \mu_{t-1,i}^{e}, \Sigma_{t-1,i}^{e}) dy_{t-1} \\ &= \sum_{i=1}^{M} w_{t,i}^{p} \phi(y_t; \mu_{t,i}^{p}, \Sigma_{t,i}^{p})\end{aligned}$$

where the weights, means, and covariance matrices are defined as

$$\begin{aligned} w^p_{t,i} &= w^e_{t-1,i} \\ \mu^p_{t,i} &= A\mu^e_{t-1,i} \\ \Sigma^p_{t,i} &= A\Sigma^e_{t-1,i}A^\top + R \end{aligned}$$

B. Measurement Step

$$\begin{aligned} bel(y_t) &= c_t p(z_t \mid y_t)\overline{bel}(y_t) \\ &= c_t\phi(z_t; h(x_t, y_t), Q)\sum_{i=1}^{M} w^p_{t,i}\phi(y_t; \mu^p_{t,i}, \Sigma^p_{t,i}) \\ &= \sum_{i=1}^{M} c_t w^p_{t,i}\phi(z_t; h(x_t, y_t), Q)\phi(y_t; \mu^p_{t,i}, \Sigma^p_{t,i}) \\ &= \sum_{i=1}^{M} c_t w^p_{t,i}\phi(z_t; \mu^z_{t,i}, \Sigma^z_{t,i})\cdot \\ &\qquad \frac{\phi(z_t; h(x_t, y_t), Q)\phi(y_t; \mu^p_{t,i}, \Sigma^p_{t,i})}{\phi(z_t; \mu^z_{t,i}, \Sigma^z_{t,i})} \\ &= \sum_{i=1}^{M} w^e_{t,i}\phi(y_t; \mu^e_{t,i}, \Sigma^e_{t,i}) \end{aligned}$$

where the weights, means, and covariance matrices are defined as

$$\begin{aligned} w^e_{t,i} &\triangleq c_t w^p_{t,i}\phi(z_t; \mu^z_{t,i}, \Sigma^z_{t,i}) \\ \mu^e_{t,i} &= h(x_t, \mu^p_{t,i}) \\ \Sigma^e_{t,i} &= (\mathbb{I} - K_t H(x_t, \mu^p_{t,i}))\Sigma^p_{t,i} \\ K_t &= \Sigma^p_{t,i}H(x_t, \mu^e_{t,i})^\top \left[H(x_t, \mu^e_{t,i})\Sigma^p_{t,i}H(x_t, \mu^e_{t,i})^\top + Q\right]^{-1} \end{aligned}$$

and $H(x_t, \mu^e_{t,i})$ is the Jacobian of $h(x_t, y_t)$ evaluated at $y_t = \mu^e_{t,i}$.

C. Measurement Distribution

$$\begin{aligned} p(z_t) = \frac{1}{c_t} &= \int p(z_t \mid y_t)\overline{bel}(y_t)dy_t \\ &= \int \phi(z_t; h(x_t, y_t), Q)\sum_{i=1}^{M} w^p_{t,i}\phi(y_t; \mu^p_{t,i}, \Sigma^p_{t,i})dy_t \\ &= \sum_{i=1}^{M} w^p_{t,i}\int \phi(z_t; h(x_t, y_t), Q)\phi(y_t; \mu^p_{t,i}, \Sigma^p_{t,i})dy_t \\ &= \sum_{i=1}^{M} w^p_{t,i}\phi(z_t; \mu^z_{t,i}, \Sigma^z_{t,i}) \end{aligned}$$

where

$$\begin{aligned} \mu^z_{t,i} &= h(x_t, \mu^p_{t,i}) \\ \Sigma^z_{t,i} &= H(x_t, \mu^p_{t,i})\Sigma^p_{t,i}H(x_t, \mu^p_{t,i})^\top + Q \end{aligned}$$

D. Measurement Distribution Over a Time Horizon

$$\begin{aligned} p(z_{1:T}) &= \int p(z_{1:T}, y_{0:T})dy_{0:T} \\ &= \int p(z_{1:T} \mid p(y_{0:T}))p(y_{0:T})dy_{0:T} \\ &= \int \prod_{t=1}^{T}\phi(z_t; h(x_t, y_t), Q)\prod_{t=1}^{T} p(y_t \mid y_{t-1}) \\ &\qquad p(y_0)dy_{0:T} \\ &= \int \prod_{t=1}^{T}\phi(z_t; h(x_t, y_t), Q)\prod_{t=1}^{T}\phi(y_t; Ay_{t-1}, R) \\ &\qquad \sum_{i=1}^{M} w^0_i\phi(y_0; \mu^0_i, \Sigma^0_i)dy_{0:T} \\ &= \sum_{i=1}^{M} w^0_i\int \phi(z_{1:T}; h(x_{1:T}, y_{1:T}), \tilde{Q})\phi(y_{1:T}; \tilde{A}y_0, \tilde{R}) \\ &\qquad \sum_{i=1}^{M} w^0_i\phi(y_0; \mu^0_i, \Sigma^0_i)dy_{0:T} \\ &= \sum_{i=1}^{M} w^z_i\phi\left(z_{1:T}; \mu^z_i, \Sigma^z_i\right) \end{aligned}$$

where

$$\begin{aligned} w^z_i &= w^0_i \\ \mu^z_i &= h(x_{1:T}, A_{1:T}\mu^0_i) \\ \Sigma^z_i &= H(x_T, \mu^0_i)\left[A_{1:T}\Sigma^0_i A^\top_{1:T} + R_{1:T}\right]H(x_T, \mu^0_i)^\top + Q_{1:T} \end{aligned}$$

VIII. Technology Transitions

The techniques developed through this thrust correspond the MAST programs mission to develop the technology to enable true autonomous operation of each member of the team.

References

[1] D. Mehrholz, L. Leushacke, W. Flury, R. Jehn, H. Klinkrad, and M. Landgraf, "Detecting, tracking and imaging space debris," *ESA Bulletin(0376-4265)*, no. 109, pp. 128–134, 2002.

[2] S. T. Acton and N. Ray, "Biomedical image analysis: tracking," *Synthesis Lectures on Image, Video, and Multimedia Processing*, vol. 2, no. 1, pp. 1–152, 2006.

[3] B. Grocholsky, "Information-theoretic control of multiple sensor platforms," Ph.D. dissertation, Citeseer, 2002.

[4] G. M. Hoffmann and C. J. Tomlin, "Mobile sensor network control using mutual information methods and particle filters," *IEEE Transactions on Automatic Control*, vol. 55, no. 1, pp. 32–47, 2010.

[5] B. J. Julian, M. Angermann, M. Schwager, and D. Rus, "A scalable information theoretic approach to distributed robot coordination," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 5187–5194.

[6] B. Charrow, "Information-Theoretic Active Perception for Multi-Robot Teams," Ph.D. dissertation, University of Pennsylvania, 2015.

[7] M. Schwager, P. Dames, D. Rus, and V. Kumar, "A multi-robot control policy for information gathering in the presence of unknown hazards," in *International Symposium on Robotics Research*, 2011, pp. 17–33.

[8] A. Rényi, "On measures of entropy and information," in *Fourth Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, 1961, pp. 547–561.

[9] E. Stump, V. Kumar, B. Grocholsky, and P. M. Shiroma, "Control for localization of targets using range-only sensors," *The International Journal of Robotics Research*, vol. 28, no. 6, pp. 743–757, 2009.
[10] N. Atanasov, J. Le Ny, K. Daniilidis, and G. Pappas, "Decentralized Active Information Acquisition: Theory and Application to Multi-Robot SLAM," in *IEEE Int. Conf. on Robotics and Automation*, 2015, pp. 4775–4782. [Online]. Available: http://dx.doi.org/10.1109/ICRA.2015.7139863
[11] D. Gu and H. Hu, "Rényi entropy based target tracking in mobile sensor networks," *IFAC Proceedings Volumes*, vol. 44, no. 1, pp. 13 558–13 563, 2011.
[12] B. Charrow, V. Kumar, and N. Michael, "Approximate representations for multi-robot control policies that maximize mutual information," *Autonomous Robots*, vol. 37, no. 4, pp. 383–400, 2014.
[13] T. M. Cover and J. A. Thomas, "Elements of information theory 2nd edition," 2006.
[14] F. Wang, T. Syeda-Mahmood, B. C. Vemuri, D. Beymer, and A. Rangarajan, "Closed-form jensen-renyi divergence for mixture of gaussians and applications to group-wise shape registration," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2009, pp. 648–655.
[15] M. Huber, *Nonlinear Gaussian Filtering: Theory, Algorithms, and Applications*. KIT Scientific Publishing, 2015, vol. 19.
[16] S. Thrun, W. Burgard, and D. Fox, "Probabilistic robotics," 2005.

2017 Final Report Metrics

Task Number: CNC-16-17.2.1

Metrics

Notes: *Please list metrics for the reporting quarter only and only Publications that are accepted and searchable on web.*

Awards and/or Recognitions this Quarter

Vijay Kumar: IEEE Robotics and Automation Society: 2017 George Saridis Leadership Award in Robotics and Automation

Refereed/Journal Papers accepted this quarter

N/A

Conference Presentations/proceedings presented this quarter

N/A

Joint Papers with ARL presented or accepted this quarter – duplication with journals/ conferences expected.

N/A

Invention Disclosures for this quarter

N/A

Patent Applications Filed this quarter

N/A

Patents Awarded this quarter

N/A

PhDs Award this quarter:

N/A

Masters Awarded this quarter

N/A

Students supported fully or partially (by name)

MS: N/A
PhD: *Shih-Ling Phuong, Santiago Paternain Olascoaga, Nikolay Atanasov, Philip M. Dames, Kelsey L. Saulnier*
Post Docs: N/A
Master Degree: N/A
PhDs Awarded: N/A

MAST CTA Member Rotations into ARL

N/A

ARL Researchers Working at Member Sites

N/A

Workshops and Seminars conducted:

N/A

MAST Quarterly Metrics Summary Table (Numbers in this table should correspond to the list above):

MAST Final REPORT METRICS TABLE (The data entered into this table should be for the reporting quarter only)	
Awards (Miscellaneous Recognitions and Awards)	1
Refereed/Journal papers	0
Conference presentations/proceedings	0
Joint Papers with ARL	0
Invention disclosures	0
Patent applications filed	0
Patents awarded	0
PhDs Award	0
Masters Awarded	0
Post Docs Supported	0
Master Degree Students Supported	0
Undergraduate Students Supported	0
CTA Member rotations into ARL	0
Visiting Researchers from ARL to the Member Sites	0

Vision-Based Cooperative Control for Aerial Vehicles

CNC-16-17.3.2

Spyridon Leonardos, Giuseppe Loianno, Kostas Daniilidis and Vijay Kumar
School of Engineering and Applied Sciences
University of Pennsylvania
Philadelphia, PA 19104, USA
{spyridon, loiannog, kostas, kumar}@seas.upenn.edu

***Abstract*— This report focuses on the vision-based state estimation and cooperative sensing component of the task. The main technical challenge is the extension of state estimation and sensing techniques to multirobot systems with minimal exchange of information and computational effort and most importantly, in a decentralized setting. In addition, a semi-distributed approach for cooperative localization and mapping with Micro Aerials Vehicles using smartphones devices is presented.**

***Index Terms*— Cooperative state estimation, distributed vision algorithms, robust Kalman filtering, data association, distributed bundle adjustment.**

I. INTRODUCTION

Collective state estimation and cooperative sensing consist fundamental problems in multirobot systems and essential for performing cooperative tasks such as formation control. In this task, we consider the case of aerial vehicles equipped with cameras and IMUs. The main challenges can be summarized as:

- Decentralization of the estimation process: to ensure that the state estimator is scalable with the number of the robots in the group, each robot should maintain only its state estimate and not the state of the entire group. As a result, the correlations between different robots estimates are unknown. In other words, the collective error covariance matrix is only partially maintained.
- Data association: matching commonly observed visual features is required for obtaining inter-robot measurements but it is a rather challenging task.
- Building a 3D map of the environment in a distributed fashion: it is challenging in the absence of a central processing node. Sensors must solve an optimization problem by communicating only with a small subset of adjacent sensors.

This report is structured as follows: Section II summarizes the robust Kalman filter approach which enables the decentralization of collective state estimation. Section III includes two methods for decentralized data association. Section IV includes the proposed approach for building a 3D map of the environment in a distributed fashion and finally, Section V summarizes a proposed approach for multiple aerial smartphone navigation.

II. ROBUST KALMAN FILTER

State estimation is one of the fundamentals problem in control theory and robotics. The most common state estimators are undoubtedly the Kalman filter [14], which is optimal in the minimum mean squared error for the case of linear systems, and its generalizations for nonlinear systems: the Extended Kalman Filter (EKF) [24] and the Uscented Kalman Filter (UKF) [11].

In multi-agent systems, the task of state estimation takes a collaborative form in the sense that it involves inter-agent measurements and constraints. Examples are cooperative localization in robotics [23] using relative pose measurements, camera network localization using epipolar constraints [25] and many more. On the one hand, a decentralized solution that scales with the number of agents is necessary. On the other hand, the state estimates become highly correlated as information flows through the network. Ignoring these correlations has grave consequences: estimates become optimistic and result in divergence of the estimator. This phenomenon is analogous to the rumor spreading in social networks.

Unknown correlations may be present in other scenarios as well. A popular simplification, that significantly reduces computations and enables the use of EKF-based estimators, is that noise sources are independent. For instance, a common assumption in vision-aided inertial navigation is independence of the image projection noises for each landmark [9], [20] although in reality, they are coupled with the motion of the camera sensor. In other cases, it might be impractical to store the entire covariance matrix due to storage limitations. For instance, in Simultaneous Localization and Mapping (SLAM) problems, there have been several approaches that decouple the sensor state estimate from the estimates of the landmark positions [13], [20] to increase the efficiency of the estimator and reduce the storage requirements.

The most popular algorithm for fusion under the presence of unknown correlations is the Covariance Intersection (CI) method which was introduced by Julier and Uhlmann [12]. In its simplest form, the Covariance Intersection algorithm is designed to fuse two random vectors whose correlation is not known by forming a convex combination of the two estimates in the information space. Covariance Intersection produces estimates that are provably consistent, in the sense

that estimated error covariance is an upper bound of the true error covariance. However, it has been observed [2], [21], [27] that Covariance Intersection produces estimates that are too conservative which may decrease the accuracy and convergence speed of the overall estimator when used as a component of an online estimator.

In our recently submitted paper [15], we propose a minimax formulation for fusion under the presence of unknown correlations. We extend our formulation to linear measurement models that depend on two random vectors with unknown cross-correlations. As an application we consider the problem of decentralized state estimation for a group of agents. The proposed estimator takes cross-correlations into account while being less conservative than the widely used Covariance Intersection. We demonstrate the superiority of the proposed method compared to Covariance Intersection with numerical examples and simulations within the specific application of decentralized state estimation using relative position measurements.

One of the most prominent applications of the proposed fusion algorithm is distributed state estimation in an EKF-based framework. However, the problem of distributed state estimation is far from new. There have been numerous approaches for EKF-based distributed state estimation and EKF-based cooperative localization. Yet, some of them require that each agent maintains the state of the entire network [23], which is impractical and does not scale with the number of agents, while others ignore correlations [19] in order to simplify the estimation process or use Covariance Intersection and variations of it despite its slow convergence.

We assume we have two random vectors x, y with expectations

$$E[x] = \overline{x}, \quad E[y] = \overline{y}. \tag{1}$$

and a linear measurement model of the form

$$z = C\overline{x} + D\overline{y} + \eta \tag{2}$$

where the measurement noise has zero-mean normal distribution, i.e.

$$\eta \sim N(0, \Sigma_\eta) \tag{3}$$

We assume that the measurement noise process is independent of the estimates. As in the classic Kalman filter derivation, we are looking for an update step of the form

$$\widetilde{x}^+ = \widetilde{x} + K(z - \widetilde{z}) \tag{4}$$

where

$$\widetilde{z} \doteq C\widetilde{x} + D\widetilde{y} \tag{5}$$

. The error of the update is given by

$$\widetilde{x}^+ - \overline{x} = (I - KC)(\widetilde{x} - \overline{x}) - KD(\widetilde{y} - \overline{y}) + K\eta \tag{6}$$

and the associated error covariance is given by

$$\begin{aligned}\widetilde{\Sigma}^+_{xx} = &\begin{bmatrix} I - KC & -KD \end{bmatrix} \begin{bmatrix} \Sigma_{xx} & \Sigma_{xy} \\ \Sigma_{xy}^T & \Sigma_{yy} \end{bmatrix} \begin{bmatrix} I - C^T K^T \\ -D^T K^T \end{bmatrix} \\ &+ K\Sigma_\eta K^T \end{aligned} \tag{7}$$

where the cross-correlation term should satisfy the following linear matrix inequality:

$$\begin{bmatrix} \Sigma_{xx} & \Sigma_{xy} \\ \Sigma_{xy}^T & \Sigma_{yy} \end{bmatrix} \succeq 0 \tag{8}$$

in order to be a valid cross-correlation. By using Schur complement, the minimax formulation is written as follows

$$\begin{aligned} \underset{K}{\text{minimize}} \quad & \sup_{\Sigma_{xy}} \quad \operatorname{tr}(\Sigma^+_{xx}) \\ \text{subject to} \quad & \Sigma_{yy}^{-1/2}\Sigma_{xy}^T\Sigma_{xx}^{-1}\Sigma_{xy}\Sigma_{yy}^{-1/2} - I \preceq 0 \end{aligned} \tag{9}$$

The numerical method we use to solve the problem at hand is an interior point method specifically designed for convex-concave games. First, we will look at the simpler case of an unconstrained convex-concave game, that is a problem of the form

$$\underset{u}{\text{minimize}} \ \underset{u}{\text{maximize}} \ f(u, v) \tag{10}$$

where the pay-off function f is convex in the first argument and concave in the second argument. The optimality conditions for differentiable convex-concave pay-off function are given by

$$\nabla_u f(u^\star, v^\star) = 0, \qquad \nabla_v f(u^\star, v^\star) = 0 \tag{11}$$

We use the infeasible start Newton method [4], outlined below, to find the optimal solution of the unconstrained problem.

Algorithm 1 Infeasible start Newton method.

given: starting points $u, v \in \mathbf{dom} f$,
tolerance $\epsilon > 0$, $\alpha \in (0, 1/2)$, $\beta \in (0, 1)$.

Repeat

1. $r(u, v) = [\nabla_u f(u, v)^T, \nabla_v f(u, v)^T]^T$
2. Compute Newton steps by solving
$$Dr(u, v)[\Delta u_{nt}, \Delta v_{nt}] = -r(u, v)$$
3. Backtracking line search on $\|r\|_2$.
   $t = 1$.
   $u_t = u + t\Delta u_{nt}$, $v_t = v + t\Delta v_{nt}$.
   While $\|r(u_t, v_t)\|_2 > (1 - \alpha t)\|r(u, v)\|_2$
      $t = \beta t$.
      $u_t = u + t\Delta u_{nt}$, $v_t = v + t\Delta v_{nt}$.
   EndWhile
4. Update: $u = u + t\Delta u_{nt}$, $v = v + t\Delta v_{nt}$.

until $\|r(u, v)\|_2 \leq \epsilon$

However, the problem at hand is slightly more complicated since it involves a linear matrix inequality. Therefore, we use the barrier method [4]. Intuitively, a sequence of unconstrained minimization problems is solved, using the last point iteration is the starting point for the next iteration.

Next we present a numerical example which sheds light on the differences between the Covariance Intersection (CI) and the proposed Robust Fusion (RF) approaches. Consider

the example of fusing two random variables with zero means and covariances given by

$$\Sigma_{xx} = \begin{bmatrix} 5 & 0 \\ 0 & 5 \end{bmatrix}, \ \Sigma_{yy} = \begin{bmatrix} 3 & 0 \\ 0 & 7 \end{bmatrix} \tag{12}$$

The corresponding error covariances obtained from Covariance Intersection and Robust Fusion are

$$\Sigma_{CI} = \begin{bmatrix} 3.79 & 0 \\ 0 & 5.79 \end{bmatrix}, \ \Sigma_{RF} = \begin{bmatrix} 3 & 0 \\ 0 & 5 \end{bmatrix}, \tag{13}$$

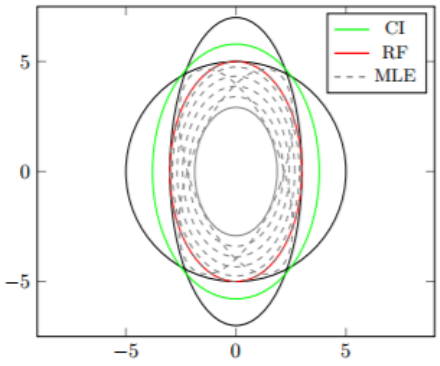


Fig. 1. Confidence ellipses: given a covariance matrix Σ we draw the set $\{x : x^T\Sigma^{-1}x = 1\}$. Initial confidence ellipse (black), Maximum likelihood Estimate (MLE) confidence ellipse (gray dashed) for various values of correlation, CI confidence ellipse (green) and RF confidence ellipse (red). The confidence ellipses obtained from MLE lie in the intersection of the two ellipsoids $\{x : x^T\Sigma_{xx}^{-1}x \leq 1\}$ and $\{x : x^T\Sigma_{yy}^{-1}x \leq 1\}$. Since the proposed approach is equivalent to MLE for some worst-case correlation, the RF confidence ellipse lies in the intersection of the two ellipsoids as well. When correlation increases, the trace of the covariance of MLE approaches the trace of Σ_{RF}. CI is not maximum likelihood for any value of correlation but produces a guaranteed upper bound on the true error covariance.

Finally, we consider an application in distributed state estimation using relative position measurements. We experiment with a group of 4 agents on the plane. If there is an edge from i to j, then agent i transmits its current state estimate and the corresponding error covariance estimate to agent j which upon receipt, takes a measurement of the relative position and updates its own state estimate and associated error covariance estimate. All agents have identical second order integrator dynamics. Only one agent is equipped with global position system (GPS). Updates of the state estimates from relative position measurements can be performed by either ignoring cross-correlations (Naive Fusion) or by one of Covariance Intersection or the proposed Robust Fusion. Each agent maintains only its one state and communicates its to each neighbors at each time instance. The individual prediction step is the same as the Kalman Filter (KF) prediction step.

We evaluate four estimators, three decentralized and one centralized: Naive Fusion (NF) which ignores correlations, Robust Fusion (RF), Covariance Intersection (CI) and Centralized Kalman Filter (CKF). The Centralized Kalman Filter (CKF) is simply a standard Kalman Filter containing all agent states. It serves as a measure of how close the decentralized estimators are to the optimal centralized estimator. The Robust Fusion based estimator significantly outperforms the Covariance Intersection based estimator which produces particularly noisy velocity estimates.

III. DISTRIBUTED CONSISTENT DATA ASSOCIATION

Multi-sensor data association has been a long standing problem in robotics and computer vision. It refers to the problem of establishing correspondences between feature points, regions, or objects observed by different sensors and serves as the basis for many high-level tasks such as localization, mapping and planning. Most of the efforts in previous works have been dedicated to improving the data association by designing new feature detectors, descriptors, and outlier removal algorithms in a pairwise setting. However, the problem setting in practice is often multi-way if the data are collected by a sensor network or a multi-robot system, and how to establish consistent data association for multiple sensors and leverage the global reasoning to improve the association has drawn increasing attention in recent years. In our recent paper [16], we propose two fully decentralized methods for consistent global data association from pairwise data associations.

A necessary condition for good data association among multiple sensors is the cycle consistency meaning that the composition of associations along a cycle of sensors should be identity. The cycle consistency is often violated in practice due to outliers and how to use such a constraint to remove the false associations has been studied in robotics, computer vision and graphics.

The cycle consistency has been explored in a variety of computer vision applications. For example, Zach et al. [28] studied the problem of multi-view reconstruction and proposed to identify unreliable geometric relations (e.g. relative poses) between views by measuring the consistencies along a number of cycles. Nguyen et al. [22] addressed the problem of finding point-wise maps among a collection of shapes and proposed to use the cycle consistency to identify the correctness of maps and replace the incorrect maps with the compositions of correct ones. Zhou et al. [29] proposed to solve the problem by rank minimization and developed a more efficient algorithm. All of these works addressed problem in a centralized setting, i.e., all pairwise measurements are available and optimized jointly. In the robotics community, numerous approaches for associating elements of two sets have been proposed and extensively used in single robot Simultaneous Localization And Mapping (SLAM) settings. More recently, multi-robot localization and mapping algorithms have been developed. In almost all of them, the data association is addressed either in a pairwise fashion by directly generalizing data association techniques from the single robot case [5], [6], [26] or by considering all associations jointly in a centralized manner [10]. Closer to our approach is the work by Aragues et al. [1] which is a decentralized method for finding inconsistencies based on cycle detection. However, the proposed inconsistency resolution algorithm comes without any optimality or guarantees.

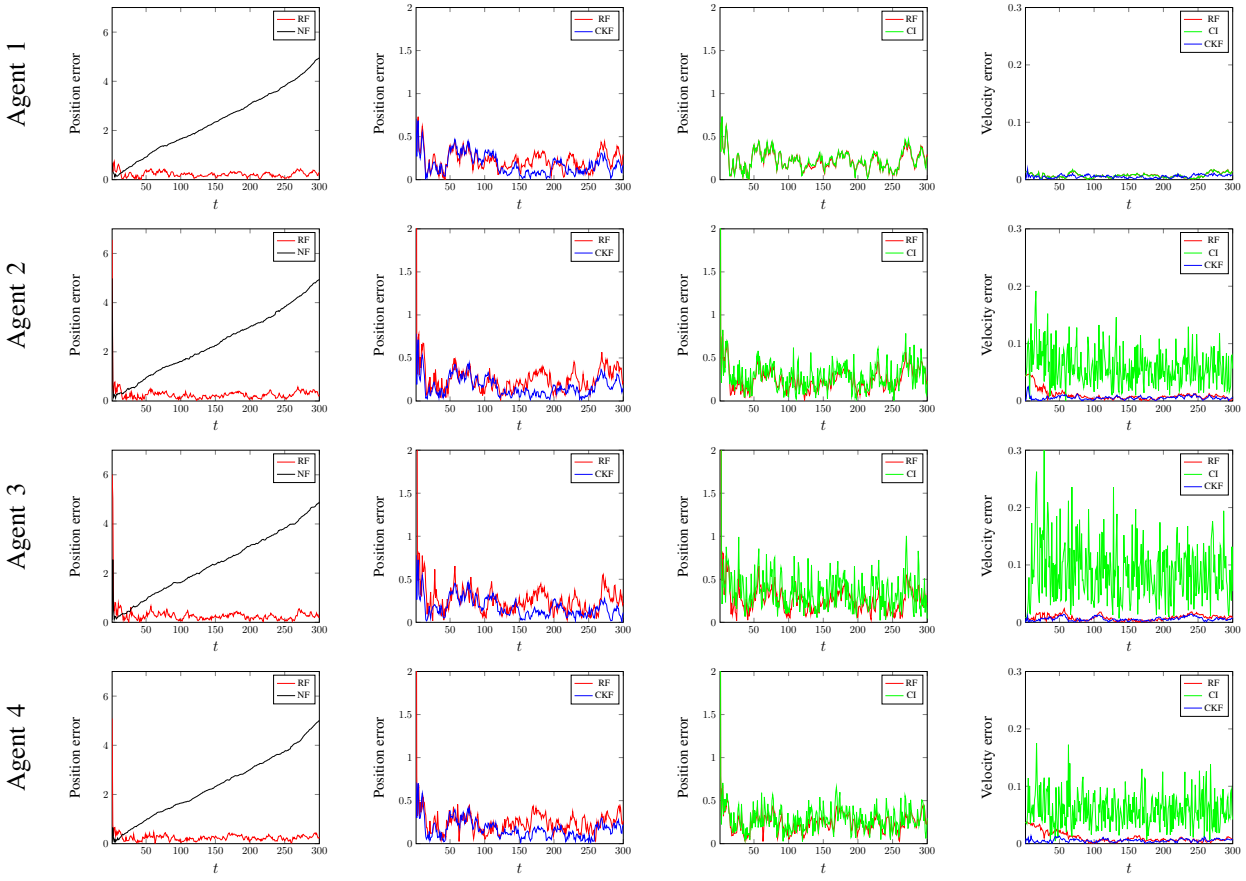


Fig. 2. Comparison of the four methods. The Naive Fusion estimator quickly diverges, whereas the Robust Fusion and Covariance Intersection do not. Clearly the Robust Fusion estimator is more accurate than the Covariance Intersection and produces less noisy estimates due to its less conservative nature. The steady state error covariance estimate of the Covariance Intersection is much larger than the actual error covariance making the estimator more susceptible to measurement noise.

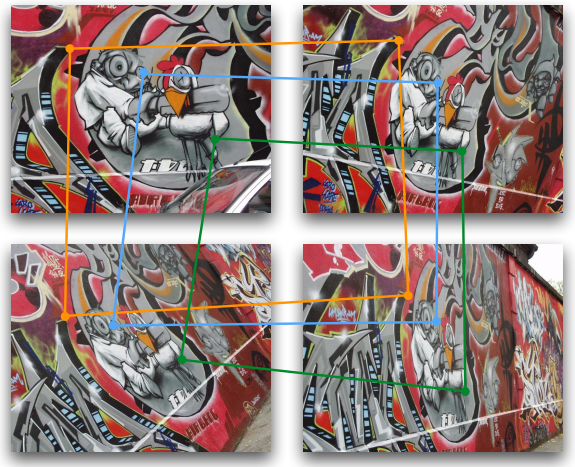


Fig. 3. Illustration of consistent data association.

More technically, the set of all permutations of n elements forms a group under composition, termed the *symmetric group*. A permutation π is represented by a n-by-n permutation matrix Π such that by

$$[\Pi]_{ij} = 1, \quad \text{if} \quad \pi(j) = i \tag{14}$$

and 0 otherwise. The simplest choice for a distance on the set of permutations is given by

$$d(\Pi_1, \Pi_2) = n - \text{tr}(\Pi_1^T \Pi_2) \tag{15}$$

which is simply the number of labels assigned differently by the two permutations.

We assume there are n sensors observing m targets. For instance, assume we have a set of n cameras observing a scene in the world described by a set of m feature points. Sensors communicate only with a subset of all sensors. Communication constraints between sensors are encoded by the sensor graph. The sensor graph is the digraph

$$G = (\mathcal{V}, \mathcal{E}), \qquad \mathcal{V} = \{1, 2, \ldots, n\} \tag{16}$$

and (i,j) is in the edge set if there is information flow from sensor i to sensor j. The pairwise association π_{ij} is defined as follows: we have that

$$\pi_{ij}(l) = k \tag{17}$$

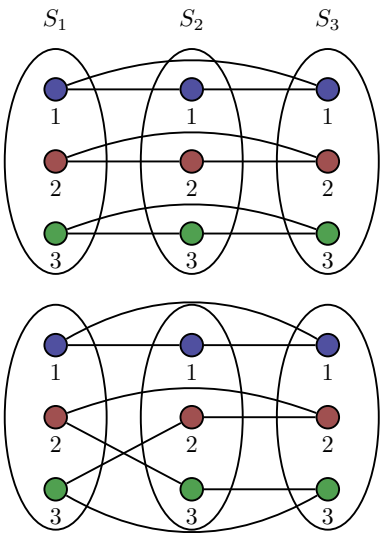


Fig. 4. Example with $n = 3$ sensors S_1, S_2, S_3 observing $m = 3$ targets. Top: consistent data association. Bottom: inconsistent data association since $\pi_{12} \circ \pi_{23}(2) = 3$ but $\pi_{31}(2) = 2$.

if the lth target in jth sensor corresponds to the kth target in ith sensor. Observe that the pairwise associations can be written as

$$\pi_{ij} = \pi_i \circ \pi_j^{-1}, \tag{18}$$

where π_i is the mapping from the labels of sensor i to some global labels, termed the "universe of features" in some works [29]. This is analogous to the absolute position in Euclidean spaces in localization problems. We denote by $\widetilde{\pi}_{ij} \in \mathfrak{S}_m$ the, possibly erroneous, estimated pairwise association between sensor i and j and by $\widetilde{\Pi}_{ij}$ the corresponding matrix representation. Related to the sensor graph is the *data association graph*

$$\mathcal{D} = (\mathcal{V}_\mathcal{D}, \mathcal{E}_\mathcal{D}, w_\mathcal{D}), \quad \mathcal{V}_\mathcal{D} = \mathcal{V} \times \{1, 2, \ldots, m\}. \tag{19}$$

There is an edge from (i, k) to (j, l) if and only if

$$(i, j) \in \mathcal{E} \qquad \text{and} \qquad [\widetilde{\Pi}_{ij}]_{kl} > 0. \tag{20}$$

First, we need a precise definition of consistency. A set of pairwise associations is consistent if

$$\widetilde{\pi}_{ij} \circ \widetilde{\pi}_{jk} = \widetilde{\pi}_{ik} \tag{21}$$

for all valid indices i, j, k. Based on the definition of consistency, the problem of consistent data association is naturally defined as follows: given pairwise associations $\{\widetilde{\pi}_{ij}\}_{(i,j)\in\mathcal{E}}$, find labels $\pi_1, \ldots, \pi_n$, such that

$$\widetilde{\pi}_{ij} = \pi_i \circ \pi_j^{-1}, \qquad \forall\ (i, j) \in \mathcal{E} \tag{22}$$

as much as possible according to some fitness criterion.

The first obstacle one encounters is the combinatorial nature of the problem at hand. Therefore, we relax the domain of problem and work with doubly stochastic matrices instead of permutation matrices. This is a natural choice since the convex hull of the set of m-by-m permutation matrices is exactly the set of m-by-m doubly stochastic matrices. The first of the two solution we propose is based upon the following update rule:

$$\Pi_i(t+1) = \frac{1}{|\mathcal{N}_i| + 1}\big(\Pi_i(t) + \sum_{j\in\mathcal{N}_i} \widetilde{\Pi}_{ij}\Pi_j(t)\big) \tag{23}$$

After the convergence of Protocol (23), we discretize the solution by solving the assignment problem which can be solved using the Hungarian algorithm . If the sensor graph contains a rooted-out branching and the pairwise associations are noiseless, then the consensus protocol with the distinguished vertex converges to a set of consistent labels up to a global permutation. However, in the general case, pairwise associations will contain noise. Therefore, we should prove convergence without assuming perfect associations. If the sensor graph contains a rooted-out branching, then the consensus protocol with a distinguished vertex asymptotically converges to a solution that depends only on the input pairwise associations and on the initial value of the labels of the distinguished vertex.

As a second method for the data association problem, we propose a decentralized implementation of a spectral method. The multi-sensor data association can be cast as the following optimization problem:

$$\underset{\Pi_1,\Pi_2,\ldots,\Pi_n\in\mathfrak{S}_m}{\text{maximize}} \sum_{(i,j)\in\mathcal{E}} \text{tr}(\Pi_i^T \widetilde{\Pi}_{ij} \Pi_j) \tag{24}$$

In the above form, the problem is computationally intractable. However, if the individual constraints are relaxed into a single collective orthogonality constraint, the relaxed problem becomes computationally tractable since it becomes a Rayleigh quotient problem:

$$\begin{aligned} \underset{\mathbf{\Pi}}{\text{maximize}} \quad & \text{tr}(\mathbf{\Pi}^T \mathbf{P} \mathbf{\Pi}) \\ \text{subject to} \quad & \mathbf{\Pi}^T \mathbf{\Pi} = m I_m \end{aligned} \tag{25}$$

where

$$[\mathbf{P}]_{ij} = \widetilde{\Pi}_{ij}. \tag{26}$$

It is a well-known fact that the solution of the above problem is given by the matrix having as columns the m leading eigenvectors of P. Numerically, the optimal solution in a centralized setting can be computed using Orthogonal Iteration [7]. Intuitively, each iteration consists of two main steps: a power step and an orthogonalization step based on QR-decomposition.

Algorithm 2 Orthogonal Iteration

Require: Input matrix $\mathbf{P}$, initial $\mathbf{\Pi}(0)$
1: **for** $t = 0, \ldots, N$ **do**
2: $\quad \mathbf{Y}(t) = \mathbf{P}\mathbf{\Pi}(t)$
3: $\quad \mathbf{Q}(t)\mathbf{R}(t) = \mathbf{Y}(t)$
4: $\quad \mathbf{\Pi}(t+1) = \mathbf{Q}(t)$
5: **end for**

The Orthogonal Iteration method is not readily decentralizable since the QR-decomposition step requires information

from the entire sensor graph. Nevertheless, observe that

$$\mathbf{R}^T\mathbf{R} = \mathbf{Y}^T\mathbf{Y} = \sum_{i=1}^{n} Y_i^T Y_i. \tag{27}$$

If the above sum can be computed in a decentralized manner, then we can recover R(t) using Cholesky factorization. Let

$$Z_i \doteq nY_i^T Y_i \tag{28}$$

$$\overline{Z} \doteq \frac{1}{n}\sum_{i=1}^{n} Z_i = \mathbf{R}^T\mathbf{R} \tag{29}$$

Based on this observation, if all sensors know n and solve a consensus problem for computing the average, then the orthogonalization step can be performed locally using Cholesky decomposition.

Algorithm 3 Distributed Orthogonal Iteration

Require: Input associations $\{\widetilde{\Pi}_{ij}\}_{(i,j)\in\mathcal{E}}$, initial $\{\Pi_i(0)\}_{i\in\mathcal{V}}$

1: **for** $t = 0, \ldots, N$ **do**
2: $Y_i(t) = \sum_{j\in\mathcal{N}_i\cup\{i\}} \widetilde{\Pi}_{ij}\Pi_j(t)$
3: Estimate $\overline{Z}_i(t)$ by consensus.
4: $R_i(t) = \text{chol}(\overline{Z}_i(t))$
5: $\Pi_i(t+1) = Y_i(t)R_i(t)^{-1}$
6: **end for**

The solution obtained by the Orthogonal Iteration is optimal up to an arbitrary m-by-m orthogonal transformation. Finally, the corrective orthogonal transformation is computed by solving an Orthogonal Procrustes problem.

First, we experiment with synthetic data to validate the accuracy of proposed methods across different problem settings. We generate instances with $m = 50$ targets and $n \in \{20, 50, 100\}$ sensors. As a performance criterion we use the accuracy of the obtained labels, which is simply defined as the percentage of correct labels. We vary the percentage of erroneous associations in each pairwise association from 10% to 90% to test the robustness of each method against the presence of outliers. We first experiment with a fully connected graph, and then we remove half of the edges. In the case of a fully connected graph, the consensus protocol achieves exact recovery of the labels for up to approximately 40% of outliers. The spectral method outperforms the consensus protocol and is able to achieve exact recovery for up to 80%-90% of outliers. By reducing either the number of edges in the sensor graph or the number of sensors, the accuracy of both methods slightly drops. As expected, the benefit of using joint matching in accuracy increases with the size of the network and the number of edges.

Next, we experiment with the CMU hotel sequence which consists of 111 frames and the CMU house sequence which consists of 101 frames. We detect corner features in the first image that remain visible across the entire sequence. The ground-truth associations are obtained by tracking the features across the entire sequence. Specifically, we used 89 points in the hotel sequence and 55 in the house sequence. Pairwise matches are computed between an image and at most 10 previous images and 10 subsequent images in the sequence, thus the sensor graph is relatively sparse. Pairwise matches are obtained by extracting SIFT descriptors and then solving the assignment problem using the Hungarian algorithm. The baseline method consists of obtaining labels by directly matching with the first image in the sequence. The ground-truth associations are extracted by tracking the features over the sequence. Both methods significantly outperform the baseline which shows that joint matching can significantly improve the overall matching performance.

TABLE I
CMU DATASETS RESULTS

Dataset	Baseline	Spectral	Consensus
House	0.716	0.861	0.887
Hotel	0.666	0.872	0.878

Next, we experiment with the Affine Covariant Regions Datasets which consist of sequences of 6 images with significant overlap but with viewpoint variability. We detect approximately 300 SIFT features [17] in each sequence that are visible across the entire sequence. The ground-truth associations are extracted using the available ground-truth homographies. Both methods significantly outperform the baseline in the two challenging sequences Graffiti and Wall. However, in the remaining sequences the input pairwise associations are already very accurate resulting in only minor improvements if any. It is worth noting that the spectral method and the consensus method have similar performance in all real sequences despite the fact that the spectral method is computationally more expensive. Therefore, in situations where the computational resources are limited, for instance swarms of UAVs with on-board cameras, the consensus algorithm would be preferable.

TABLE II
GRAFFITI DATASETS RESULTS

Dataset	Baseline	Spectral	Consensus
Graffiti	0.544	0.639	0.647
Wall	0.570	0.646	0.658
Trees	0.912	0.902	0.909
Leuven	0.962	0.972	0.979
Bikes	0.986	0.987	0.987
UBC	0.996	0.996	0.996

IV. DISTRIBUTED BUNDLE ADJUSTMENT

Given a set of images depicting a number of 3D points from different viewpoints, bundle adjustment can be defined as the problem of jointly estimating the 3D locations of scene points, the parameters of the relative motion and possibly the optical characteristics of the camera sensors, according to an optimality criterion involving the corresponding image projections of all points.

In a centralized setting, bundle adjustment is commonly formulated as the following optimization problem:

$$\underset{\{R_i,T_i,X^j\}}{\text{minimize}} \quad \sum_{i,j} \|\widetilde{x}_i^j - \pi(R_i^T(X^j - T_i))\|_2^2 \tag{30}$$

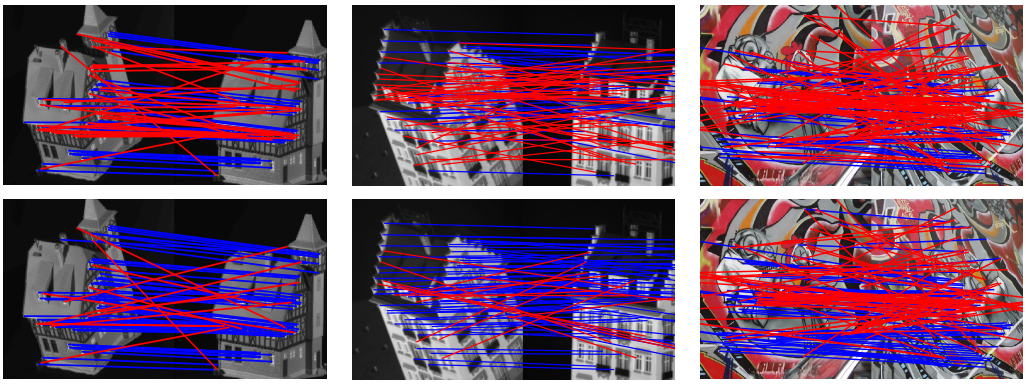


Fig. 5. Top: initial matches. Blue color corresponds to inlier matches whereas red color corresponds to outlier matches. Bottom: matches after the consensus algorithm.

where $(R_i, T_i) \in SE(3)$ denotes the pose of the ith visual sensor, $X^j \in \mathbb{R}^3$ denotes the 3D position of the jth observed feature point, $\widetilde{x}_i^j$ denotes the noisy projection of the jth point in the ith view, if it is visible and finally π is the perspective function.

In a decentralized setting, that is, in the absence of central processing node that maintains the estimated 3D map, each visual sensor i maintains local estimates of the visible 3D points. Let X_i^j denote the local estimate of the jth point by the ith visual sensor. We formulate the problem of distributed bundle adjustment as the following equality-constrained optimization problem:

$$\begin{aligned} \underset{\{R_i,T_i,x_i^j,X_i^j,\lambda_i^j,X^j\}}{\text{minimize}} \quad & \sum_{i,j} \|\widetilde{x}_i^j - x_i^j\|_2^2 \\ \text{subject to} \quad & X_1^j = \cdots = X_m^j = X^j \\ & \begin{bmatrix} x_i^j \\ 1 \end{bmatrix} = \frac{1}{\lambda_i^j} R_i^T (X_i^j - T_i) \end{aligned} \tag{31}$$

Intuitively, we minimize the reprojection error subject to consensus constraints and projection equations constraints. Since the above optimization problem has only equality constraints, we solve by alternating minimization on its augmented Lagrangian.

Next we present some results in the figure that follows, where both the input images and the reconstructed point clouds and camera poses are included. We achieve accurate localization of the camera sensors and reconstruction of the point cloud.

V. MULTIPLE AERIAL SMARTPHONE NAVIGATION

The architecture enabling flight of multiple aerial platforms powered by smartphones is mostly distributed. Each vehicle is able to perform state estimation, planning and control tasks. The only role of the ground station is to solve the goal assignment problem and to collect data from each platform such to build a cooperative map of the environment. Multiple vehicles are able to plan safe trajectories avoiding inter-robot collisions, optimizing at the same time a given task and concurrently building in a cooperative manner a 3-D map of the environment. The sensing, sensor fusion, control, and planning are all done on an off-the-shelf Samsung Galaxy S5 smartphone using just the single camera and IMU available on the phone. The work allows any consumer with multiple smartphones to autonomously drive a swarm of multiple vehicles without GPS, by downloading an app, and have the swarm cooperatively map a 3-D environment.

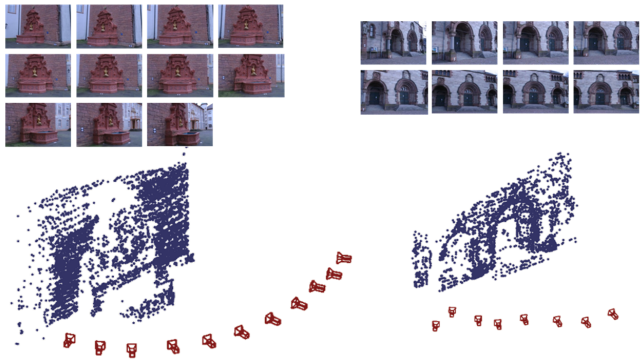


Fig. 6. Results of the proposed approach for distributed bundle adjustment.

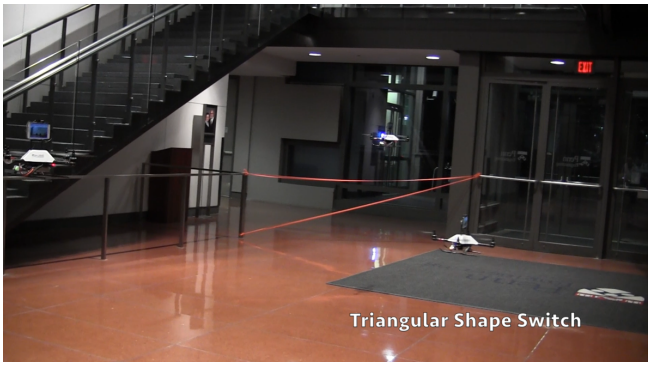


Fig. 7. Formation navigation with multiple vehicles.

VI. CONCLUSIONS AND FUTURE WORK

During the last year, we have worked on fundamental aspects of the vision-based distributed state estimation prob-

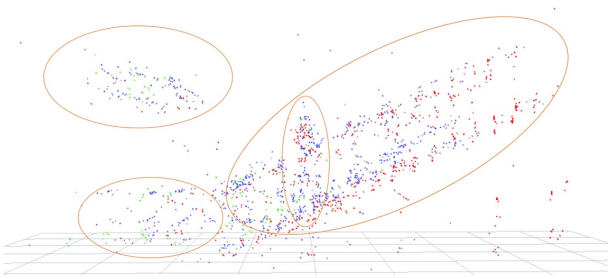


Fig. 8. Cooperative reconstruction from multiple vehicles.

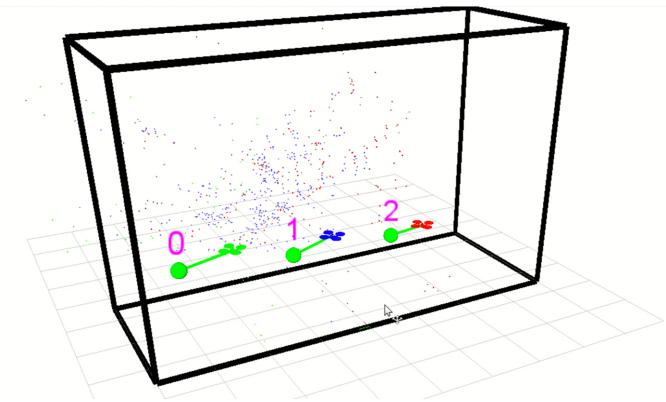


Fig. 9. User interface with goal selections and reconstructions.

lem. Robust Kalman Filtering, can be used in any collaborative localization framework and additional applications are worth investigating. The ultimate goal is to integrate all components in a single distributed collective odometry framework that achieves robust and reliable vision-based state estimation for groups of robots.

REFERENCES

[1] Rosario Aragues, Eduardo Montijano, and Carlos Sagues. Consistent data association in multi-robot systems with limited communications. In *Robotics: Science and Systems*, pages 97–104, 2011.

[2] Pablo O Arambel, Constantino Rago, and Raman K Mehra. Covariance intersection algorithm for distributed spacecraft state estimation. In *American Control Conference, 2001. Proceedings of the 2001*, volume 6, pages 4398–4403. IEEE, 2001.

[3] Mica Arie-Nachimson, Shahar Z Kovalsky, Ira Kemelmacher-Shlizerman, Amit Singer, and Ronen Basri. Global motion estimation from point matches. In *3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT), 2012 Second International Conference on*, pages 81–88. IEEE, 2012.

[4] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[5] Alexander Cunningham, Kai M Wurm, Wolfram Burgard, and Frank Dellaert. Fully distributed scalable smoothing and mapping with robust multi-robot data association. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 1093–1100. IEEE, 2012.

[6] Dieter Fox, Jonathan Ko, Kurt Konolige, Benson Limketkai, Dirk Schulz, and Benjamin Stewart. Distributed multirobot exploration and mapping. *Proceedings of the IEEE*, 94(7):1325–1339, 2006.

[7] Gene H Golub and Charles F Van Loan. *Matrix computations*, volume 3. JHU Press, 2012.

[8] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004.

[9] J. Hernandez, K. Tsotsos, and S. Soatto. Observability, identifiability and sensitivity of vision-aided inertial navigation. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2319–2325, May 2015.

[10] Vadim Indelman, Erik Nelson, Nathan Michael, and Frank Dellaert. Multi-robot pose graph localization and data association from unknown initial relative poses via expectation maximization. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 593–600. IEEE, 2014.

[11] Simon J Julier and Jeffrey K Uhlmann. New extension of the kalman filter to nonlinear systems. In *AeroSense'97*, pages 182–193. International Society for Optics and Photonics, 1997.

[12] Simon J Julier and Jeffrey K Uhlmann. A non-divergent estimation algorithm in the presence of unknown correlations. In *In Proceedings of the American Control Conference*. Citeseer, 1997.

[13] Simon J Julier and Jeffrey K Uhlmann. Using covariance intersection for slam. *Robotics and Autonomous Systems*, 55(1):3–20, 2007.

[14] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Journal of basic Engineering*, 82(1):35–45, 1960.

[15] Spyridon Leonardos and Kostas Daniilidis. A game-theoretic approach to robust fusion and kalman filtering under unknown correlations. *arXiv preprint arXiv:1610.01045*, 2016.

[16] Spyridon Leonardos, Xiaowei Zhou, and Kostas Daniilidis. Distributed consistent data association. *arXiv preprint arXiv:1609.07015*, 2016.

[17] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004.

[18] Yi Ma, Stefano Soatto, Jana Kosecka, and S. Shankar Sastry. *An Invitation to 3-D Vision: From Images to Geometric Models*. SpringerVerlag, 2003.

[19] Agostino Martinelli. Improving the precision on multi robot localization by using a series of filters hierarchically distributed. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1053–1058. IEEE, 2007.

[20] Anastasios I Mourikis and Stergios I Roumeliotis. A multi-state constraint kalman filter for vision-aided inertial navigation. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pages 3565–3572. IEEE, 2007.

[21] Esha D Nerurkar and Stergios I Roumeliotis. Power-slam: A linear-complexity, consistent algorithm for slam. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 636–643. IEEE, 2007.

[22] Andy Nguyen, Mirela Ben-Chen, Katarzyna Welnicka, Yinyu Ye, and Leonidas Guibas. An optimization approach to improving collections of shape maps. *Computer Graphics Forum*, 30(5):1481–1491, 2011.

[23] Stergios I Roumeliotis and George A Bekey. Distributed multi-robot localization. *IEEE Transactions on Robotics and Automation*, 18(5):781–795, 2002.

[24] Harold Wayne Sorenson. *Kalman filtering: theory and application*. IEEE, 1985.

[25] Roberto Tron and Rene Vidal. Distributed 3-d localization of camera sensor networks from 2-d image measurements. *Automatic Control, IEEE Transactions on*, 59(12):3325–3340, 2014.

[26] Stefan B Williams, Gamini Dissanayake, and Hugh Durrant-Whyte. Towards multi-vehicle simultaneous localisation and mapping. In *Robotics and Automation, 2002. Proceedings. ICRA'02. IEEE International Conference on*, volume 3, pages 2743–2748. IEEE, 2002.

[27] Xun Xu and Shahriar Negahdaripour. Application of extended covariance intersection principle for mosaic-based optical positioning and navigation of underwater vehicles. In *Robotics and Automation, 2001. Proceedings 2001 ICRA. IEEE International Conference on*, volume 3, pages 2759–2766. IEEE, 2001.

[28] Christopher Zach, Manfred Klopschitz, and Manfred Pollefeys. Disambiguating visual relations using loop constraints. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2010.

[29] Xiaowei Zhou, Menglong Zhu, and Kostas Daniilidis. Multi-image matching via fast alternating minimization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4032–4040, 2015.

Rapid Deployment of Heterogeneous Robots: Abstractions, Algorithms, and Experimentation

CNC-16-3-3

Stefano Carpin, Jose Susa Rincon
School of Engineering
University of California, Merced
Merced, CA, USA

Vijay Kumar, Kostas Karydis
Dept. of Mechanical Engineering
University of Pennsylvania
Philadelphia, PA, USA

Herbert G. Tanner
Dept. of Mechanical Engineering
University of Delaware
Newark, DE, USA

***Abstract*—We consider a problem dubbed *rapid deployment,* whereby a team of robots is tasked with providing situational awareness to the warfighter by collecting relevant information in an area to be entered by the soldiers at a later time. The emphasis of this task is on *operational tempo*, i.e., robots are required to complete their information-gathering task while performing at a speed comparable with humans'. This problem is cast as a resource allocation problem, i.e., we determine how many robots are needed to complete the task within a given temporal deadline and with a predefined level of confidence. At the same time, we provide also control policies for the robot(s) to obtain the desired outcome. Different scenarios are considered with regard to the number of robots available, the number of locations to be visited, and data collection capabilities. To bridge the gap between higher level planning tools and low-level controllers, suitable abstractions and models are introduced. Finally, in an effort close the loop between theory and practice, all components developed in this task are being combined together to experimentally verify their soundness.**

***Index Terms*—Situational Awareness, Multi-robot Systems, Resource Allocation.**

I. Introduction

The operational vision of the MAST program encompasses missions with a short time horizon and *high operational tempo*, e.g., the last hundred meters separating the soldiers from a site to be entered and inspected. In such a tactical scenario, time is of essence. Accordingly, this task considers the problem of deploying multiple robots in a given environment with the objective of increasing situational awareness for the warfighter tasked with clearing a space of interest. The team of robots enters this environment concurrently or shortly before the warfighter, with the objective to visit an assigned list of relevant locations (e.g., rooms, corridors, etc.) and report any potential threat detected. There is therefore an inherent temporal coupling between robots and humans, because the robotic team must conclude its task before warfighters commence theirs. This coupling forces the robotic team to operate at a time scale comparable with human operational tempo. Deploying at operational tempo inevitably increases uncertainty and failure rates since robots operate at the limit of their capabilities. Therefore, it is necessary to quantitatively model and predict how the performance of the team is affected when operating under increasingly more demanding conditions.

The parameters of our problem are as follows: we have numerous robotic platforms, including aerial and ground vehicles. Few of them are reliable, predictable, and expensive. These are robots equipped with state-of-the-art sensors, communication, and computational devices. The remaining ones (a much larger number) are expendable, small-scale (e.g., palm-sized) platforms with limited on-board sensing, communication, and computational power. The dynamics of these small-scale platforms are not precisely known, likely variable, and it is expected that they will be subjected to noise and environmental disturbances during their deployment. These platforms, moreover, are expendable. The environment in which these robots operate is assumed to be known or at least partially known. The questions we address in this task are the following.

1. What are the levels of actuation and control authority, coordination, sensing fidelity, team composition, communication bandwidth, prior information, and hardware redundancy that would guarantee timely deployment with acceptable and certifiable failure rate? For the problem at hand, deployment time and success probability are considered adjustable mission parameters.
2. What meaningful trade-offs between these general system resources can we identify through principled analysis and design methods? Currently, there exist no formal procedures to address questions of this nature. At the heart of the issue lies a resource allocation and optimization problem, the solution of which is required to strike a balance between risk quantification and mitigation, performance guarantees, and robustness to asset failure. In this specific context, such problems have not yet been framed and formalized.

In this task we have, for the first time, studied these problems in a unified way, spanning the whole spectrum from modeling to validation in the real world.

The task is solved combining three different technical contributions. First, we developed a novel framework to model in a principled way the uncertain dynamics of robot platforms considered in the MAST project. In particular, following a data-driven approach, stochastic motion primitives are generated with the property of matching experimental data. Next, these stochastic primitives are used as the action space for a model based on Constrained Markov Decision Processes (CMDP), whereby an agent subject to uncertain dynamics is tasked with achieving a certain goal while being subject to various costs that need to be wither optimized or bounded. Finally, the CMDP framework is used to formulate and solve various planning problems relevant for the MAST program.

II. Related Work

There exist a significant body of scientific literature pertaining to the use of multiple robots for cooperative collection of information. This is a common theme in the context of deployment and coverage in mobile sensor networks. A significant portion of the work on deployment is based on heuristic or rule-based behavioral approaches [1,2]. These strategies are intuitive and easily interpretable but do not allow any formal quantification of performance in terms of convergence to a desired configuration or time to completion. More formal approaches to deployment and coverage [3,4,5] offer predictable and provably convergent steady state behavior, and have been shown to be relatively robust to external disturbances and noise due to the use of feedback. The emphasis in these approaches is on convergence, which is guaranteed asymptotically; however these asymptotic bounds are typically overly conservative to make these algorithms relevant to applications with short time horizon deadlines. In addition, the majority of work on mobile sensor networks deployment makes simplifying assumptions on the dynamics of the agents, which in a high operational-tempo framework and for the size factor of MAST platforms may not always be valid: in high speeds, inertia effects and kinematic constraints matter, and at the small scale noise and model parameter uncertainty may not be ignored. Finally, there has been some work on operating small teams of heterogeneous robotic platforms for the purposes of search and coverage [6,7,8], network connectivity maintenance [9], and cooperative mapping [10], but mission completion time has been considered sporadically, and only as some metric of performance, rather than an explicit problem specifications and constraint. Meeting deadlines, on the other hand, has been partially studied as a problem in a multi-agent system setting [11,12], but not in the context of deployment in 2D or 3D real-world environments. There is therefore a gap in our understanding and our knowledge base when it comes to coordinating collections of heterogeneous, dynamically agile robotic sensor platforms that are to operate under time deadlines and in the face of uncertainty.

This project leverages ONR funding, as well as MAST funding through the task CNC 14-4.2 which led over the past years to the early results on trajectory planning on large teams of robots. Models for MAV flight were developed with previous funding under MCE 14-4.5. Preliminary results regarding decision-making using constrained Markov Decision Processes were developed with supplemental funding SUPP-13-6-CNC.

III. Approach

The problems considered in this task require a coherent, unified solution to various intertwined problems. Some of these problems are low-level control problems, whereas others deal with high-level sequential decision-making problems. In particular it is necessary to simultaneously consider:

1. The unavoidable uncertainty in platform dynamics;
2. Suitable abstractions to map continuous time/space dynamics into discrete models that can be efficiently processed and analytically characterized;
3. Planning algorithms for single and multi-agent systems under a variety of boundary conditions.

In the following we illustrate how each of these problems has been tackled.

A. Uncertainty In Platform Dynamics

The class of vehicles we consider is affected by significant process noise and it is unrealistic to expect that controllers will be capable of accurately tracking a given trajectory. It is therefore of interest to develop suitable models capable of capturing these uncertainties. To this end, we model uncertainty both in the mathematical representation of platform dynamics as well as in their mobility behavior in the form of low-dimensional continuous-time/space stochastic differential equations. This approach build upon preliminary results presented in [13], for the case of an *octoroach* robot (see Figure 1), but can be generalized to other MAST relevant platforms, too, including wheeled vehicles, as well as (micro) aerial vehicles.

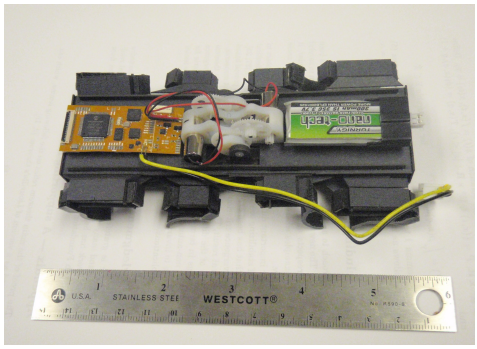


Figure 1: Octoroach robot

Keeping in mind that these models need to be incorporated into discretized representations suitable for sequential stochastic decision-making, it is of interest to determine the conditions under which a discrete model of a stochastic

process can be related through a simulation relation to the stochastic process. The methodology combines simulation relations with randomized algorithms to produce data-driven probabilistic discrete abstractions for stochastic processes lacking a concrete model. It does so in a way that the abstract model can enjoy formal fidelity guarantees with respect to the physical processes it describes. The data-driven methodology works as follows. Let $W=\{w_1, w_2, \ldots, w_I\}$ denote a sequence of I experimentally collected state observations from a dynamical system of interest. For example, each w_i could represent the sequence of planar position and orientation attained by a ground robot as it navigates for a predetermined amount of time T. Note that thanks to the generality of the presented method, w_i could also represent the position and orientation of a quadrotor or some other comparable platform, i.e., the methodology is not restricted to just one class of vehicles. Let M be a model for the system parametrized by a vector of l parameters collectively indicated as E. As we vary E, a family of models $M(E)$ is obtained, and each model is associated with an output $out(M(E))$. The identification problem, then, is to determine the vector E^* associated with the model $M(E^*)$ generating the output closest to W with respect to some appropriate metric. Indicating with $\tilde{w}$ the sample mean, this identification problem can then be solved as the following least squares optimization problem:

$$E^* = \underset{E}{\operatorname{argmin}} \sum_{t=1}^{T} \left\| out(M(E))_t - \tilde{w}(t) \right\|^2$$

The solution to this problem determines the vector E^* best fitting W in average, but does not adequately capture its variability. To capture such variability, E^* is used as the mean of a random vector whose higher order statistics are estimated from W as well. This step is called *stochastic model extension problem*, and cannot be solved by simply fitting higher order statistics as for the mean, because such approach would not provide any probabilistic guarantee about its accuracy. The guarantee is introduced as follows. For the sample set, we define a confidence level γ, i.e., the desired probability that the model matches the sample set. Next, we define a function g, indicating whether a given parameter vector E' matches the sample with the given confidence level γ. This function *geometrically* matches $out(M(E'))$ to the sample set and either accepts or rejects E'. Starting from the mean E^* identified with the least squares problem described above, values for the covariance matrix of the model are iteratively increased until the matching test succeeds, i.e., until the trajectories generated by the model match the experimental data according to the desired level of confidence γ. This process, described in greater detail in [14,15].

Once the identification problem is solved, the model can used for generation purposes, i.e., it can be used to simulate trajectories that are guaranteed to be consistent with the experimental dataset W used for training.

The process described in this subsection is repeated for different W sets, each of them representing a different *motion primitive* like for example "move straight ahead", or "turn to the right", etc. This finite set of motion primitives will then feed into the abstract representation described in the next subsection, and then finally into the planning algorithms for single and multirobot systems. Obviously, if multiple heterogeneous vehicles are considered, the method can be repeated for each vehicle, and a different set of maneuvers can be considered for each of them. Results showing this methodology are shown in the Section IV.

B. Discrete Abstractions

A suitable discretization method is needed to connect the continuous, low-level stochastic models with the high-level planners presented in the next section. Among the various options available, we opt for a representation based on Constrained Markov Decision Processes (CMDP). This approach has various advantages. First, it allows studying decision-making problems where the outcome of actions is uncertain, and this is consistent with our assumption of operating with vehicles with uncertain dynamics. Moreover, it naturally associates multiple costs to each action. In our abstraction, each action is one of the stochastic primitives characterized with the method presented in the previous subsection. Accordingly, every primitive can be characterized by various costs. The rich literature in CMDP provides formulations for multi-cost optimization problems whereby one cost is minimized (say the probability of failure), while others are bounded in expectation (e.g., the time needed to complete the assigned task). Initial efforts for using CMDP models for the rapid deployment problem were presented in [16]. A CMDP consists of the following elements:

- A discrete state space X with n states.
- For each state $x \in X$, a finite set $A(x)$ representing the set of actions that can be applied when the system is in state x. It is also convenient to consider the state/space set $K=\{(x,a) : x \in X, a \in A(x)\}$.
- For each couple of states $x,y \in X$, and every action $a \in A(x)$, the *one-step* transition probability P_{xy}^a, i.e., the probability of going from state x to state y when action a is executed.
- a primary cost function $c:K \rightarrow R$. $c(x,a)$ is the cost incurred when executing action $a \in A(x)$ while in state x.
- L additional cost functions $d_i:K \rightarrow R$, representing L additional costs incurred when executing an action from a state.

The connection between the CMDP model and the stochastic representations considered in the previous section is as follows. Each primitive is mapped into one action a that can be executed from one of the states. States represent the position of a robot in the environment, and therefore P_{xy}^a is the probability of going from position x to position y when primitive a is executed. Every time an action is executed, $L+1$ costs are incurred, i.e., c and the L costs d_i. Costs relevant for the rapid deployment problem are, for example, the time to execute a maneuver, or the probability that a maneuver will be

successfully completed. The behavior of a system modeled with a CMDP is defined introducing a policy, i.e., a function $\pi: X \rightarrow A$ specifying the action to take in every state. It is well known that functions of this type, i.e., deterministic and Markovian lead to optimal formulations in most cases of practical interest [17]. To complete the formulation of the CMDP problem, it is necessary to introduce suitable cost models. Among the different costs proposed in literature, we opt for the *total cost*, i.e., the cumulative undiscounted cost incurred during missions whose length is a priori unknown. The total cost for the primary function is:

$$c(\pi) = \sum_{t=0}^{\infty} E_{\pi}[c(X_t, A_t)]$$

where the expectation is taken with respect to the policy π followed, while X_t and A_t are the state and action at time t.. Note that for the above expression to exist, it is necessary for the sum to be finite, and this can be assumed under mild circumstances. In essence it is sufficient to assume that after a finite amount of time the robot either completes its task, or it stops working (for example because it runs out of power). The reader is referred to [16] for technical details. Similar costs can be assumed for the costs d_i. One final technical note is necessary for the cost $c(\pi)$. In general, the cost is a function not only of the policy but also of the initial mass distribution over the states, normally indicated as β. Therefore, costs in the following will be indicated as $c(\pi,\beta)$.

C. Planning

The planning problems we consider build upon the discrete abstractions we presented in the previous subsection. The basic problem is as follows.

> Given a CMDP and L bound $D_1,\ldots,D_L$, solve the following constrained optimization problem:
>
> $$\pi^* \in \operatorname{argmin} c(\pi,\beta)$$
>
> $$s.t.\ \ d_i(\pi,\beta) \le D_i \quad 1 \le i \le L$$
>
> where the policy π is taken from the set of markovian (i.e., memory less policies).

With this formulation we can solve a variety of planning problems related to rapid deployment. For example, it is possible to set the primary cost c to be the probability of failure, and consider just one additional cost d_1 to be equal to the time to complete the task. The problem, in this case therefore asks to determine the policy minimizing the failure probability (i.e., maximizing the success probability) while bounding the expected time to complete the task (where the bound is a temporal deadline D_1).

Note that this problem formulates the policy for a single robotic agent. Based on this results obtained when solving this problem, we can therefore ask the following questions:

- Given N identical robot, and a temporal deadline T, what is the probability of successfully completing the task? This is an *analysis* question, i.e., it asks to predict the performance for a pre-assigned team configuration.
- Given a temporal deadline T and a desired success probability P^*, how many robots do we need to complete the task with probability at least P^* within time T? This is *design* question, i.e., it allows to decide how to allocate resources, i.e., how many robots should be in the team.

The previous questions are formulated for the case where all robots are identical, however it can be easily extended to consider the heterogeneous case where robots with different capabilities are considered. The optimization problem we just formulated can be solved with an appropriate linear program [18]. Let $\rho(x,a)$ be a set of optimization variables, one for each element in the state/space set K. Then the optimization problem is solvable if and only if the following linear program is feasible:

$$\min_{\rho} \sum_{(x,a)\in K} \rho(x,a)c(x,a)$$

$$s.t. \sum_{(x,a)\in K} \rho(x,a)d_i(x,a) \le D_i \quad 1 \le i \le L$$

$$\sum_{y\in X} \sum_{a\in A(y)} \rho(x,a)(\delta_x - P_{xy}^a) = \beta(x) \quad \forall x \in X$$

$$\rho(x,a) \ge 0 \quad \forall (x,a) \in K$$

From the solution of the linear program a policy can extracted to drive one robot towards its target according to the parameters set in the program. As a cross result, the solution to the linear program provides also the value of the objective function, and in our case this is the failure probability.

1) Swarm case

In the swarm case we consider the scenario where the number of robots is much larger than the number of locations to be visited, but all robots are equal. Under these conditions robustness is achieved through redundancy, i.e., the

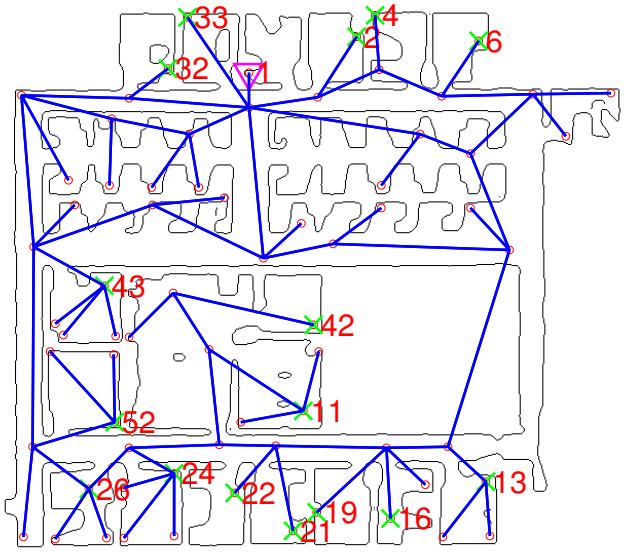


Figure 2: Test environment for the deployment problem in the swarm case

probability of successfully completing the task increases because many robots are assigned to each location to be visited. The optimization problem presented above provides the control strategy for a single robot, as well as the associated success probability. By solving this problem for each assigned location, it is possible to determine how many robots should be sent to each place of interest. This target assignment problem is an integer optimization problem and therefore hard to solve, but approximate solutions are possible (see [16] for details). Figure 2 shows an example of the environment considered.

Each location with a number is a place to be reached by one of the robots. Location number 1, marked with a pink triangle identifies the start location, i.e., the location where all robots start in the beginning. Blue edges display the connectivity between the different locations in the environment. An edge between two vertices in the graph means that it is possible to move between the two vertices, but this does not necessarily mean that a straight path is followed. In term of the CMDP representation, vertices in the graph are associated with states in the state space X, whereas and edge between two vertices, means the availability of one of more actions (maneuvers) to move between the two locations. Each maneuver is characterized by the time it takes to execute, and the probability of success. The task is completed when each of the locations marked by a number is reached by at least one robot. Therefore it is important not only to compute the policy for a single robot (for every location), but to also determine how many robots should be sent to each location.

2) Constrained Case

We next considered a version of the problem where the number of robots is smaller than the number of locations to observe. This problem is somehow the dual of the previous one, i.e., whereas in the previous case we assigned multiple robots to the same location, we now assign multiple locations to the same robot. Moreover, we introduce a sensor model featuring errors in the form of missed detections. This is modeled as follows. Let t_i be a target location to be observed. To t_i we associate a *visibility set* $V(t_i)$, i.e., a set of locations from which t_i can be seen. For simplicity we assume that visibility sets are disjoint, i.e., if the robot is positioned in a location $x \in V(t_i)$, it can only observe t_i. However, observation is not error free, i.e., even if the robot is located in one of the vertices in $V(t_i)$ and it makes an observation towards t_i, it is possible to incur in a missed detection error, i.e., the element to observe may remain unseen. This error rate is modeled by the parameter $P(x,o)$, i.e., the probability to correctly make an observation of $V(t_i)$ when in $x \in V(t_i)$. Figure 3 illustrates this idea. Target locations are indicated with thick circles, whereas their visibility regions are shown with thin circles in the same color.

One of the difficulties associated with this problem is that a policy visiting the locations cannot be Markovian, i.e., memory less, because otherwise there is the risk of revisiting multiple times the same location. This problem can be circumvented by expanding the state with a limited memory describing which target locations have been seen already.

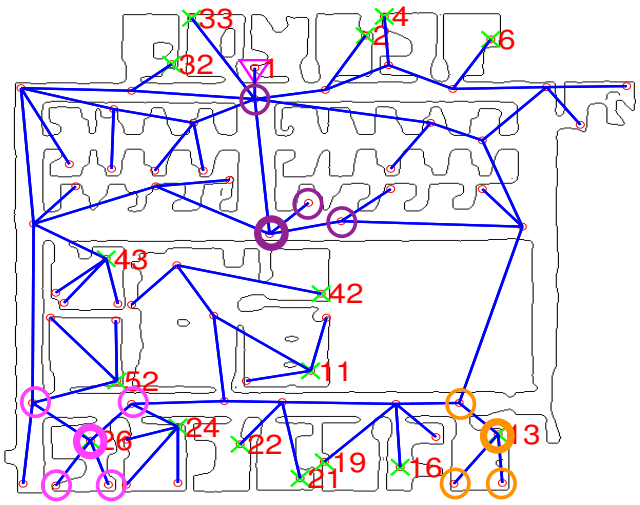


Figure 3: environment model for the scarce resources case.

However, this approach would cause a growth in the state-space by a factor of 2^m where m is the number of targets. This is somehow expected because this problem is related to the well-known traveling salesman problem (TSP). We avoided this problem by stipulating that targets are visited in a prescribed sequence without leaving to the planner the task to determine the best sequence. In this way the state space grows by a factor of m rather than 2^m. The problem to determine a suitable sequence is solved as follows. First an approximate solution for the problem is determined using a state-of-the-art solver like Concorde. This tour is then split into segments, each of them providing a candidate sequence. For each sequence, the problem of determining the optimal plan can also be cast as a CMDP instance. In this case the primary objective is to visit as many target locations as possible, subject to constraints on the failure probability P_f and the time to completion T_f. While the time to completion is usually provided as an input parameter, the failure probability is normally unknown, but it can be determined through an iterative binary search. Once a policy has been generated for all sequences it remains to be determined how to assign them to the robots for execution. It is possible to show that this problem can be solved using a greedy strategy and this leads to an approximation algorithm with factor $1-1/e$. In this case the greedy approach consists in casting the problem as a special instance of the well-known static weapon assignment problem [19]. More details about this problem are found in our recently published paper [20].

3) Other planning problems

In addition to the planning problems described above, we studied two other variations that are relevant for future research directions. In [21] we considered the problem of planning paths for a team of robots engaged in tasks like coverage, persistent monitoring, and surveillance applications. This problem, related to the constrained case considered above, asks to determine paths for all robots such that every target is visible from at least one path. Figure 4 illustrates an instance of this

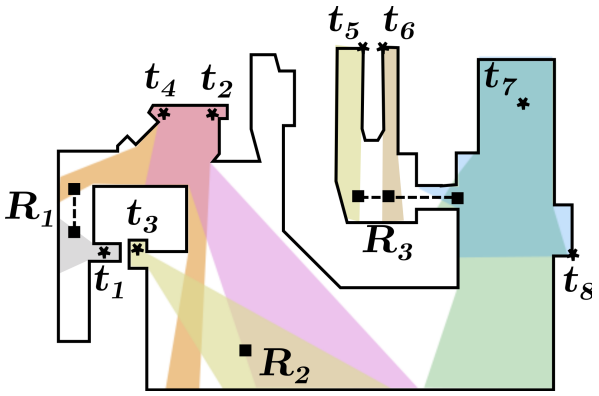


Figure 4: an example of continuous monitoring. Each of the target locations marked with a star must be seen by at least one of the robots making observations at locations marked with a square.

problem. Each of the locations marked with a star and indicated as t_i must be observed by at least one robots. Paths traveled by robots are marked by dashed paths (---) and observation points are marked with squares. The figure shows a case with eight locations to be observed. They can be observed by three robots (R_1, R_2, R_3) with R_1 making two observations along its path, R_3 making three observations along its path, and R_2 making one observation while standing still. The interesting question to ask is how many robots are necessary to solve this task, their paths, and the observation points. The main result covered in [21] is that the problem in its general form is NP-hard but in case the paths are restricted to be on pre-assigned curves the problem can be efficiently solved. In particular, we if the curves satisfy a property called *chain-visibility* then the optimal solution can be found, whereas if the boundary of the environment is polygonal, it is possible to find a constant factor approximation.

Finally, in [22] we considered the problem of *risk-aversion* since it is extremely relevant for time-sensitive tasks like rapid deployment. Risk-aversion deals with the problem of deviations from the expected behavior. The solution to the CMDP planning problems described above is in expectation only, i.e., it minimizes the expected cost and it satisfies the bounds in expectation only. While this is desirable over many iterations, there is no guarantee that one may occasionally experience an outlier realization, i.e., a realization whose cost is far away from the expectation. In tasks like rapid deployment predictability is important, and therefore one would prefer a policy generating realizations with possibly higher expectation but lower variance, so that the expected behavior is more representative of actual realizations. Combining expectation and variance into the objective function is known to be computationally hard. In [22], for the first time, we unveiled a method that achieves this objective for the case of Markov Decision Processes using the total cost criterion we adopted for the rapid deployment task. The mathematical details are rather complex and the reader is referred to [22] for more details. This work represents just the first step towards a promising direction, but work remains to be done to reduce its computational complexity (polynomial, but with an extremely large number of variables), and to extend it to the CMDP case including multiple costs.

IV. Results

Various tests have been performed to determine the soundness of the techniques developed during this task. In this section we briefly recap some of the most compelling results, and we refer the reader to the cited papers for more details.

Stochastic primitives were extensively studied in [13,14,15]. The main result is displayed in Figure 5. The figure shows 250 observed paths for the Octoroach for each of the three primitives "move forward", "move forward left," and "move forward right." The figure confirms that the robot is subject to significant uncertainty, as shown by the distribution of the end locations. The dashed curves mark the means and boundaries of the tolerance regions for each primitive for a confidence level of 0.65. The figure confirms the accuracy of the stochastic models associated with the primitives.

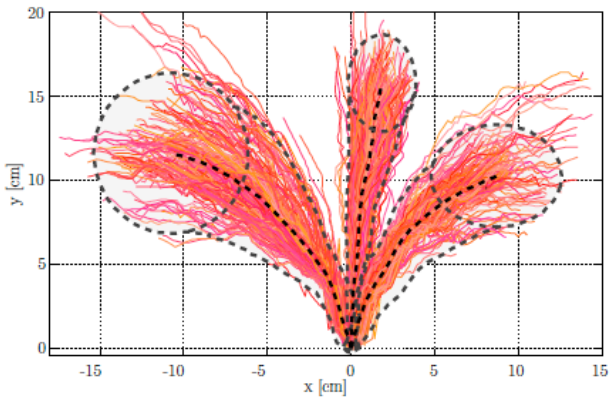


Figure 5: traces of 250 realizations for three maneuvers and their associated confidence levels.

The planning problems for the swarm case were theoretically studied [16,23], while more practical evaluations were presented in [24,25]. Figure 6 shows one of the more insightful results, namely how the parameters "team size" (*K,* on the horizontal axis), "temporal deadline" (time threshold, horizontal axis), and "success probability" (vertical axis) influence each other. The chart was derived for the environment shown in Figure 2. For any given choice of two of the parameters, the third one can be derived from the curve.
In Figure 7 we show the results for the deployment problem with scarce resources, i.e., when the number of targets exceeds the number of robots. The chart is referred to the environment presented in Figure 3 and features a map with 70 vertices and 10 locations of interest to be observed (marked as thick circles in the map). The chart shows how the probability of successfully completing the task varies as a function of the

time allocated to complete the task. Six curves are shown, corresponding to the cases where the number of robots varies from one to six. In all cases, the number of targets is ten.

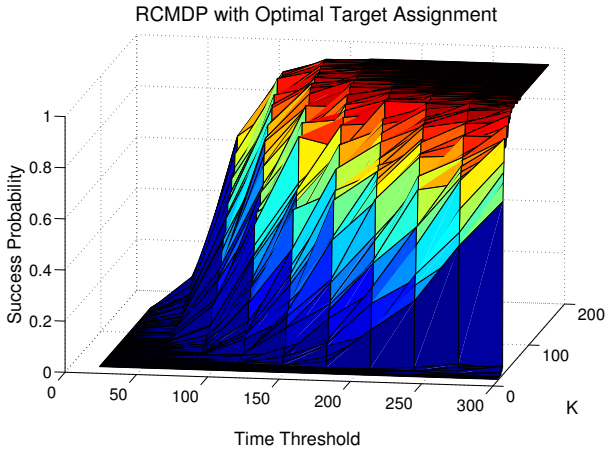


Figure 6: interplay between team size, temporal deadline and probability of success.

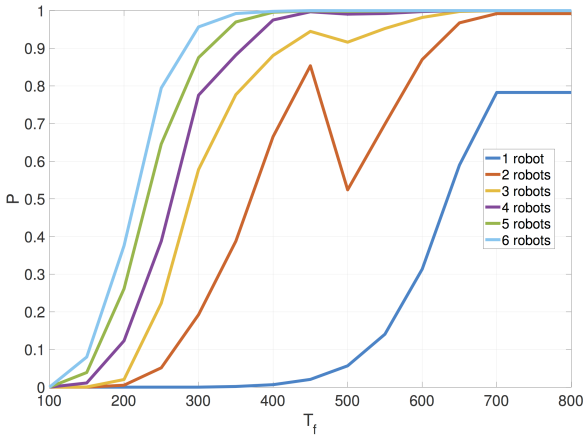


Figure 7: success rates for the case where the number of targets (10) exceeds the number of robots.

Finally, in figure 8 we plot the trend for the sum of the time to solve all the linear programs associated with a specific temporal deadline. This chart is relevant to assess the scalability of our proposed solution. The curve shows mean and standard deviation of the trends for different temporal deadlines. The small standard deviation outlines that this is largely independent from the parameter T_f and shows a quadratic growth.

V. Conclusions and Future Work

In this project we have developed the theoretical foundations for the rapid deployment problem, a task highly

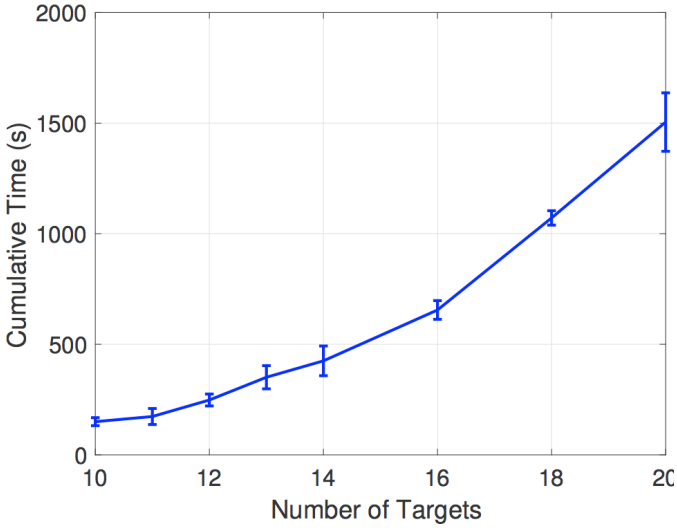


Figure 8: time spent to solve the problem as a function of the number of targets.

relevant for soldiers operating in tactical scenarios. The objective is to develop teams of robots working at operational tempo to provide situational awareness to the warfighter. The main focus of this task has been on the theoretical side, i.e., understanding the right modeling and abstractions to determine how the different parameters interact and their tradeoffs. We have shown how the three components we developed can be connected in a unified framework, and efforts to perform more extensive validations on MAST related hardware are underway at the time of writing. The various components have also been combined in an end-to-end framework integrating all of them together (a final publication describing these results is being finalized at the time of writing).

There are numerous venues for future research in this domain. First, more extensive validation is necessary, in particular on high-velocity vehicles, like micro UAVs, to understand the practical impact of design choices, like the number of maneuvers, or the discretization resolution. Initial validation efforts indicate that their impact is significant, but no general-purpose rules have been so far identified. Another important aspect to be consider concerns state observability, or localization. Using an approach based on Markov Decision Processes implies the development of a feedback policy mapping states into actions. In the rapid deployment scenario, this implies knowing the location of the robot, at least in an approximate sense. Removing this assumption comes at a significant cost, either in computational terms (because this implies entering the domain of partially observable MDPs) or in terms of performance because it will take much longer to the robot to complete the assigned task. Consequently, it appears that the development of novel technologies for indoor localization continues to be one of the most pressing issues.

In conclusion, in this task we have studied a new class of problems highly relevant to the MAST vision and with great practical importance. The theoretical foundations of rapid deployment have been studied and the proposed framework

appears to be appropriate and sound. Transitioning these findings into a ready to deploy system, however, is a problem deserving additional work and system integration efforts.

REFERENCES

[1] R. Arkin, Behavior Based Robotics, MIT Press, 1998.

[2] A. Howard, M.J. Mataric, and G.S. Sukhatme. Mobile Sensor Network Deployment using Potential Fields: A Distributed, Scalable Solution to the Area Coverage Problem. In Proceedings of the 6th International Symposium on Distributed Autonomous Robotics Systems, 2002.

[3] F. Bullo, J. Cortes, and S. Martinez. Distributed control of robotic networks: A Mathematical Approach to Motion Coordination Algorithms, Princeton University Press, 2009.

[4] D.V. Dimarogonas, S.G. Loizou, K.J. Kyriakopoulos, and M.M. Zavlanos. A feedback stabilization and collision avoidance scheme for multiple independent non-point agents . Automatica, 42(2):229-243, 2006.

[5] H.G. Tanner, A. Boddu. Multi-agent navigation functions revisited. IEEE Transactions on Robotics 28(2):1346-1359, 2012.

[6] B. Grocholsky, J. Keller, V. Kumar, and G. Pappas. Cooperative Air and Ground Surveillance. IEEE Robotics & Automation Magazine, 13(3):16-25, 2006.

[7] L.C.A. Pimenta, V. Kumar, R.C. Mesquita, and G.A.S. Pereira. Sensing and Coverage for a Network of Heterogeneous Robots, In Proceedings of the 2008 IEEE Conference on Decision and Control, 3947-3952.

[8] H.G Tanner. Switched UAV-UGV Cooperation Scheme for Target Detection. In Proceedings of the 2007 IEEE International Conference on Robotics and Automation, 3457-3462.

[9] R.A Cortez, R. Fierro, and J. Wood. Connectivity Maintenance of a Heterogeneous Sensor Network. In Proceedings of the 10th International Symposium on Distributed Autonomous Robotic Systems, 2012.

[10] N. Michael, S. Shen, K. Mohta, Y. Mulgaonkar, V. Kumar, K. Nagatani, Y. Okada, S. Kiribayashi, K. Otake, K. Yohsida, K. Ohno, E. Takeuchi, and S. Tadakoro. Collaborative mapping of an earthquake-damaged building via ground and aerial robots. Journal of Field Robotics 29(5):832-841, 2012.

[11] F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S.L. Smith. Dynamic Vehicle Routing for Robotic Systems. Proceedings of the IEEE, 99(9):1482-1504, 2011.

[12] V. Srivastava, R. Carli, C. Langbort, and F. Bullo. Attention allocation for decision making queues. Automatica 50(2):378-388, 2014.

[13] K. Karydis, I. Poulakakis, and H. G. Tanner. Probabilistic Validation of a Stochastic Kinematic Model for an Eight-legged Robot. In Proceedings of the 2013 IEEE International Conference on Robotics and Automation, 2562-2567.

[14] K. Karydis. A data-driven hierarchical framework for planning, navigation and control of uncertain systems: applications to miniature legged robots. PhD thesis, University of Delaware, 2015.

[15] K.Karydis, I. Poulakakis, J. Sun, H.G. Tanner. Probabilistic valid stochastic extensions of deterministic models for systems with uncertinaty. The International Journal of Robotics Research 34(10):1278-1295, 2015.

[16] Y.-L. Chow, M. Pavone, B.M. Sadler, S. Carpin. Trading safety versus performance: rapid deployment of robotic swarms with robust performance constraints. In *ASME Journal of Dynamic Systems, Measurement and Control*, 137(3):031005, 2015

[17] D. Bertsekas. Dynamic Programming and Optimal Control. Athena Scientific, 2005.

[18] E. Altman. Constrained Markov Decision Processes. Chapman & Hall/CRC.

[19] P.M. Pardalos, L.S. Pitsoulis. Nonlinear assignment problems. Kluwer academic publishing, 2000.

[20] J.L. Susa Rincon, P. Tokekar, V. Kumar, S. Carpin. Rapid Deployment of Mobile Robots under Temporal, Performance, Perception and Resource, Constraints. IEEE Robotics and Automation Letters, 2(4):2016-2023, 2017.

[21] P.Tokekar, V. Kumar. Visibility-based persistent monitoring with robot teams. Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, 3387-3394, 2015.

[22] S. Carpin, Y. Chow, M. Pavone. "Risk aversion in finite Markov Decision Processes using total cost criteria and average value at risk". Proceedings of the IEEE International Conference on Robotics and Automation, 335-342, 2016.

[23] S. Carpin, M. Pavone, B. Sadler. Rapid multirobot deployment with time constraints. Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, 1147-1154, 2014.

[24] A. Purohit, P. Zhang, B. Sadler, S. Carpin. Deployment of Swarms of Micro-Aerial Vehicles: from Theory to Practice, Proceedings of the IEEE International Conference on Robotics and Automation, 5408-5413, 2014.

[25] X. Chen, A. Purohit, C.R. Dominguez, S. Carpin, P. Zhang. DrunkWalk: Collaborative and Adaptive Planning for Navigation of Micro-Aerial Sensor Swarms. Proceedings of ACM SenSys, 295-308, 2015.

Inspection and Monitoring of Critical Infrastructure Resources

BPP-SUPP-16-17-4-CNC

Tolga Özaslan, Giuseppe Loianno
James Keller, Camillo J. Taylor
Vijay Kumar
GRASP - University of Pennsylvania
Philadelphia, USA

Jennifer Wozencraft
USACE
Kiln / Nashville, USA

***Abstract*— In this paper, we consider the state estimation and local mapping of a Micro Aerial Vehicle (MAV) inside a generalized-cylindrical tunnel using a 3D lidar and an Inertial Measurement Unit (IMU). The uniaxial and symmetrical geometry of the environment poses unique challenges both in pose estimation and mapping. Our motivation is to automate inspection of penstocks with the least operator intervention. Thus we propose a local mapper that captures the piecewise parametric geometry of a tunnel using a 3D lidar. We use a constrained Unscented Kalman Filter (UKF) on a Riemannian manifold for temporal estimation of both the robot pose and the local map. The measurement model is a constrained optimizer that respects the geometry of a tunnel. The proposed method easily adapts to changes in the tunnel geometry during traversal through bends and sections with different diameter. We present experimental results obtained in Center Hill Dam, TN.**

I. Introduction

Multi-rotor Micro Aerial Vehicles (MAV) have been the platform of choice by the robotics community in the last decade due to their manufacturability, agility and diverse payload options [1] Some prominent scenarios cut out for MAVs are inspection, air delivery, surveillance, search and rescue, real estate, entertainment, photography. Inspection tasks are an especially perfect fit for MAVs due to there being hard to reach spots inside mines, penstocks, bridges, ship vessels and various industrial infrastructure. Simultaneous Localization and Mapping (SLAM) [2], [3], non-linear control [4], path planning and navigation [5] problems have been extensively studied on these platforms. MAVs optimized for higher payload capacity, yet lightweight and small enough for indoor applications, allowed researchers to equip them with different types of sensors such as color and RGB-D cameras, 2D/3D lidars, IMU and GPS [6], [7]. Unlike the earlier studies on SLAM which focus on inferring the world state from a single type of sensor [3], [8], the emergence of MAVs that can move in the 3D space freely equipped with different sensor modalities, urged the robotics community to develop more sophisticated algorithms [7], [9], [10].

In this work we focus on inferring the parametric representation of a *piecewise-smooth-generalized-cylinder (PSGC)* shaped tunnel [11] (Sec. III-A). Although we present results from a penstock inspection scenario, there are many places we see PSGC structure such as mine shafts, caves and sinkholes, highway tunnels, building corridors, to name a few. Fig. 1 depicts several variations of PSGC environments where an MAV can navigate with the proposed method.

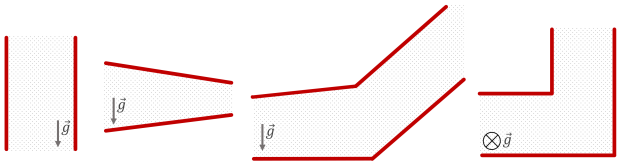


Fig. 1: We propose a state estimator and local mapper for MAV navigation inside piecewise-smooth-generalized-cylinder (PSGC) shaped environments several variations of which are illustrated. There are many places we see the PSGC structure such as mine shafts, caves and sinkholes, highway tunnels, building corridors, dam penstocks.

Penstocks are pitch-dark tunnels with diameters ranging between 5 to 20 meters and lengths between 70 to 250 meters with no geometric features other than change in bending profiles (Fig. 2). Our claim is that *uniaxial*, *symmetric* and *featureless* tunnels can be mapped with better accuracy and robustness if represented as a parametric PSGC compared to the common choice of raw pointcloud representation. Given the smoothness of the environment, we can impose roto-translational consistency constraints between the segment of a PSGC. We attain this objective through a UKF [12] estimating pose and map on $\mathbb{R} \times \mathbb{S}^2 \times \mathbb{S}^2$, a constrained non-linear optimization on $\mathbb{S}^2 \times \mathbb{S}^2$ for segment measurements and cubic Bézier interpolation for data association. We exploit spherical data analysis tools from *directional statistical literature* for noise filtering and constrained optimization on $\mathbb{S}^2$ [13].

In our previous works we used a single 2D lidar rectified with a mirror setup to reflect some of the rays to the ceiling and floor [14], and two 2D lidars one of them sweeping the ground and the other facing forward [7]. However, with only 2D slices from a 3D environment, the map and state inference problems get infeasible unless the map is known *a priori*. In our recent work [15], we used the same 3D lidar as in this work without imposing smoothness constraints on the PSGC segments. However, during our previous experiments, we experienced that when the circularity assumption fails, lack of additional constraints on the local map may fail the

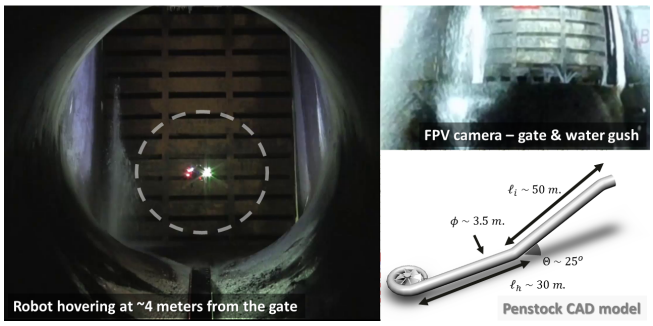


Fig. 2: *(Left)* Custom-designed hexrotor platform hovering $\sim$ 4 meters from the gate in shared-control mode. The platform is equipped with a Pixhawk autopilot, VLP-16 3D lidar, i7 NUC PC and 6S 6000mAh LiPo battery. *(Top-right)* FPV camera snapshot showing the gate and the water gush. *(Bottom-right)* CAD model of a penstock at Center Hill Dam, TN.

estimator and crash the platform, which may cause severe safety problems. Hence, an MAV should be deployed for service in such confined and wet tunnel settings only after the underlying estimator is proven to be robust. For this reason we believe that the proposed method offers a unique advantange for an MAV to survive in such environments.

Contributions of this work are three-fold. *Firstly*, we model the uniaxial, symmetric tunnel as a parametric, deformable *piecewise-smooth-generalized-cylinder* (PSGC). This inherently imposes constraints on the measurement model and avoid divergence. *Secondly*, we use Watson distribution, a statistical tool that can model axially symmetric distributions on $\mathbb{S}^p$, which is not exploited by the robotics community. We use this tool for outlier elimination on $\mathbb{S}^2$ and preserve consistency between map segments especially when the circularity assumption does not hold. *Thirdly*, we integrate sensory information for both pose and map estimation in a UKF running on a manifold. Finally, to our knowledge, this is the only MAV system that works in PSGC environments.

II. Related Work

After more than a decade of research on MAVs, these machines became the foremost robotics platforms when real-life applications are of concern. [16] is one of the interesting field studies which uses an MAV to inspect ship vessels. This study proposes an automated solution for visual inspection of internal and external parts of a ship hull using a 2D lidar, IMU and a downward facing camera. However, this work has strong assumptions on the planarity of the floor for proper working of visual odometry hence may fail when flying on curved surfaces.

In an industrial boiler inspection scenario, [17] uses a quadrotor equipped with an IMU and a stereo rig tightly coupled to estimate 6 degrees of freedom (DOF) robot state. With the on-board power-LEDs, the platform does not require any external illumination required for the vision-based estimator. However, the researchers rely on existence of significant texture on the boiler walls, which is not the case inside a penstock. Penstocks are usually coated with tar to prevent rusting but tar offers very weak texture even when it is cracked or peeled off. In a similar study, [18] inspects mine shafts with an MAV controlled manually to evaluate its performance under challenging conditions. The authors post-process the data collected with a 2D lidar and a stereo camera rig to reconstruct the vertical mine shaft. This methodology is applicable as long as the human pilot has clear view of the platform.

3D mapping of cluttered environments and state estimation of MAVs is often studied using either a 3D lidar or a 2D lidar mounted on a rotating base. [19] uses the latter hardware to map both indoor and outdoor environments with a feature-based pointcloud registration method. They claim their estimator experiences very low drift and is not demanding in terms of CPU power despite the complexity induced by integrating each 2D range reading at different times due to the hardware setup. This method is, however, not applicable in our case due to lack of geometric features inside a tunnel. In [20], the authors propose a MAV system that relies on an on-board RGB-D sensor and an IMU to estimate the robot state and map its surroundings. The proposed system is proven to work in only indoor settings due to the limited sensing capability of the RGB-D sensor. This work also demonstrates path planning and navigation using the 3D occupancy grid map previously generated by the robot.

III. Preliminaries

This study presents autonomous robust navigation of an MAV inside a featureless PSGC shaped tunnel-like environment. The robot state and the map are estimated by a UKF running on a manifold. We formulate the segment measurement function as a constrained optimization problem that respects the symmetric and smooth geometry of the environment. The optimizer fits a chain of cylindrical segments defined with their *position*, *axis* and *radius* using an on-board 3D lidar (Sec. IV-G). Consistency among the segments is satisfied using directional statistical analysis tools [13] (Sec. IV-F). We achieve temporal smoothness through a UKF running on a manifold (Sec. IV-D). In the subsequent sections, we give details on the environment, platform, notation and terminology before we present the approach in details.

A. Environment

Penstocks exhibit unique geometric properties encountered rarely in other indoor settings. They represent the characteristics of a *generalized cylinder (GC)*. [11] defines a GC as a generalization of an ordinary cylinder where the center axis can be any continuous spine and its radius may change along the axis. We write this as $GC := \{\mathbf{a}(s), r(s)\}$ where $\mathbf{a}(s)$ a vector-valued function that defines the center axis and the $r(s)$ is the radius both parametrized with $s \in \mathbb{R}$. We generalize this definition to PSGC . A PSGC is concatenation of multiple GC's where each GC is defined with its *own* axis and radius functions while preserving zeroth and first order

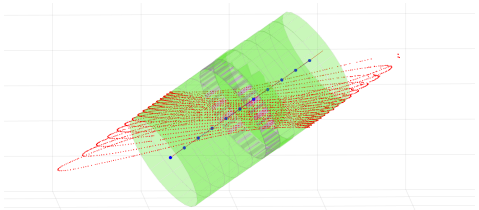


Fig. 3: A sample pointcloud data from the inclined section of a penstock in Center Hill Dam, TN. The on-board 3D lidar can only see a narrow section due to its narrow field of view. This theoretically limits the information that can be extracted for state estimation.

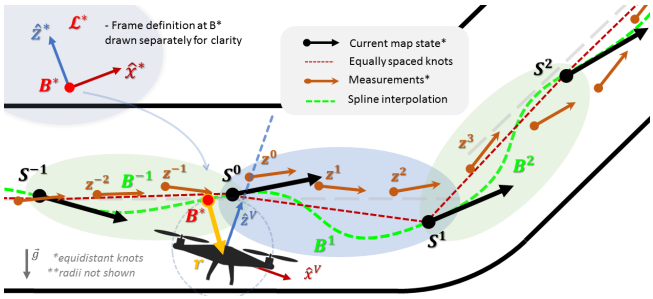


Fig. 4: An illustration of map representation, robot state, lidar measurements and data association. The black markers (knots) represent the map segments with their positions and orientations. These are used to interpolate intermediate points of the tunnel axis with a cubic Bézier polyspline (dashed green line). The orange markers represent the tunnel as given by the measurement model. These are compared against the polyspline for data association. Elliptical shades show the data association between the map segments and the measurements. Instantaneous map origin and its corresponding frame are denoted at $\mathcal{B}^*$ and $\mathcal{L}^*$. Robot state is also shown as a vector pointing from $\mathcal{B}^*$ towards the robot.

continuity with its adjacent GC's. We use cubic Bézier curves to satisfy the continuity requirements (Sec. III-C.3).

Uniaxial, symmetric tunnels lack geometric cues necessary for range-based localization. CAD model of a typical penstock is depicted in Fig. 2. The reader can observe from this CAD model that the only regions that allow for range-based 6 DOF pose estimation are the ends and junctions where sections with different axis orientations meet. Due to the 30 degrees of narrow field of view (FOV) of the lidar (Sec. III-B), this problem becomes extremely severe when the robot is flying along the inclined section (Fig. 3). The diameter of a typical penstock ranges between 5 to 15 meters and its length extends for between 70-250 meters. At the upper end of a penstock is a gate by observing which the robot can infer completion of an autonomous traversal.

B. MAV Platform

The challenges and specifications of the task require the use of a custom designed MAV platform. We use DJI F550 frame[1] with extended arms that allow using DJI E600 propulsion system[2] to obtain a high weight to thrust ratio of ~ 2. The total weight of the platform is approximately 4.5 kg when equipped with a Velodyne Puck LITE [3], Pixhawk autopilot [4], 5th generation Intel i7, 6S 5000mAh 50C LiPo battery and custom-designed power distribution board. Although we do not use in this specific scenario, the robot also carries 4 Chameleon3 1.3MP color cameras[5] and high-power Cree LEDs[6] for on-board illumination (Fig. 2). The total flight time is more than 12 minutes depending on the aggression level.

C. Notation & Terminology

1) Frame Definitions: We assume that the IMU ($\mathcal{I}$), 3D lidar ($\mathcal{V}$) and body ($\mathcal{B}$) frames are coincident due to the proper and precise placement of the sensors. The other two frames are map ($\mathcal{M}$) and world ($\mathcal{W}$) frames that differ only by yaw. We often explicitly denote the frame at which a variable is represented at as $\mathbf{x}^{\mathcal{F}}$. Vectors with no explicit frame definition should be interpreted in $\mathcal{V}$. A rotation from $\mathcal{A}$ to $\mathcal{B}$ is written as ${}^{\mathcal{B}}\mathbf{R}_{\mathcal{A}}$. Finally we denote the origin of a frame as $\mathcal{O}^{\mathcal{F}}$.

2) The Map & The Segments: We model the map, $\mathcal{M}$, as a *deformable* PSGC consisting of segments, $\{\mathcal{S}^i\}$, with fixed inter-segment Euclidean distance (Fig. 4). Segments are indexed by integers and the one closest to the robot is $\mathcal{S}^0$. A segment, $\mathcal{S}^i$, is defined with its position, local tunnel axis and its radius which respectively writes $\mathcal{S}^i := \{\mathbf{k}^i, \hat{\mathbf{t}}^i, \rho^i\}$ where $\mathbf{k} \in \mathbb{R}^3$, $\hat{\mathbf{t}} \in \mathbb{S}^2$ and $\rho \in \mathbb{R}$. Each $\mathcal{S}^i$ has its own right-handed local frame written as $\mathcal{L}^i := \{\hat{\mathbf{x}}^i, \hat{\mathbf{y}}^i, \hat{\mathbf{z}}^i\}$ which is uniquely defined by the following three relations : $\hat{\mathbf{x}}^i = \hat{\mathbf{t}}^i$, $\hat{\mathbf{y}}^i \parallel (\hat{\mathbf{t}}^i \times \hat{\mathbf{g}})$ and $\hat{\mathbf{z}}^i \cdot \hat{\mathbf{g}} < 0$ where $\hat{\mathbf{g}}$ is the gravity vector direction. The frame origin, $\mathcal{O}^i$, is the segment position, $\mathbf{k}^i$. The segments *after* and *before* the ith segment, $\mathcal{S}^i$, are denoted with the short-hand notation of $\mathcal{S}^{i^+}$ and $\mathcal{S}^{i^-}$ for $i > 0$ and vice versa for $i < 0$ (Fig. 4). The total number of forward and backward segments is denoted as N^+ and N^- excluding $\mathcal{S}^0$. We constrain the distance between adjacent segments to be constant, ℓ^i, where i is the anterior segment index.

3) Cubic Spline Interpolation: Since penstocks bend smoothly, we can assume that the points between two consecutive segments can be interpolated using a cubic Bézier spline (we will write only *Bézier* for brevity). The Bézier formed by $\mathcal{S}^{i^-}$ and $\mathcal{S}^i$ is denoted as $\mathcal{B}^i$.

A Bézier is a linear combination of 3rd order Bernstein basis polynomials with the weights being the so called *control points*. We define $\mathcal{B}^i$ with four control points the first and the last being the segment centers, $\mathbf{k}^{i^-}$ and $\mathbf{k}^i$. The other two control points are defined by projecting $\mathbf{k}^{i^-}$ and $\mathbf{k}^i$ along the segment axes, $\hat{\mathbf{t}}^{i^-}$ and $\hat{\mathbf{t}}^i$ respectively, onto the plane perpendicular to, and passing through the center of the line connecting the two segment centers (Fig. 5 shows

[1]https://www.dji.com/flame-wheel-arf
[2]https://www.dji.com/e600
[3]http://velodynelidar.com/vlp-16-lite.html
[4]https://pixhawk.org/modules/pixhawk
[5]https://www.ptgrey.com/chameleon3-usb3-vision-cameras
[6]http://www.cree.com/xlamp/xhp50_2

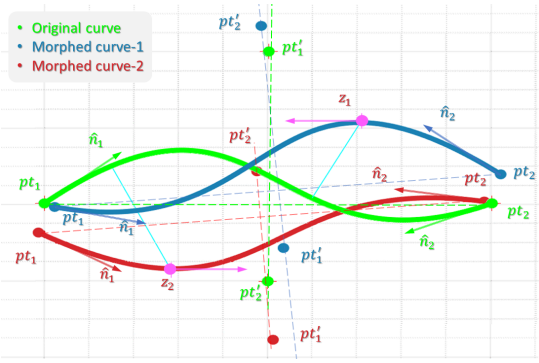


Fig. 5: This figure shows two examples of morphing a cubic Bézier spline so as to pass through a point with a specific tangent. The blue and red curves are morphed Bézier splines that pass through the purple points with illustrated tangents. End points are labeled as pt_1 and pt_2. The intermediate control points, pt_1' and pt_2', are derived from tangents, $\hat{n}_1$ and $\hat{n}_2$, given at the two end points. The morphing operation is defined as a constrained optimization problem with certain penalties on end point updates, the volume of the convex hull formed by the control and points with $\|pt_1 - pt_2\|$ kept constant. See Sec. IV-B for details.

the 2D case). Points along the curve are parametrized with $t \in [0, 1]$ where $\mathcal{B}(t=0) = \mathbf{k}^{i^-}$ and $\mathcal{B}(t=1) = \mathbf{k}^i$.

Bézier interpolation allows for defining local reference at any point along the PSGC map, between segments in a similar way as in Sec. III-C.2. The frame origin writes as $\mathcal{O}^{\mathcal{B}^i(t)}$ and the triad as $\mathcal{L}^{\mathcal{B}^i(t)} := \{\hat{\mathbf{x}}^{\mathcal{B}^i(t)}, \hat{\mathbf{y}}^{\mathcal{B}^i(t)}, \hat{\mathbf{z}}^{\mathcal{B}^i(t)}\}$. The tangent to tunnel axis is found as $\hat{\mathcal{B}}^{i'}(t) = {}^{\mathcal{B}^{i'}(t)}/_{\|\mathcal{B}^{i'}(t)\|}$ where $\mathcal{B}^{i'}(t) = {}^{d\mathcal{B}^i(t)}/_{dt}$.

We also define the distance of a point, $\mathbf{p}$, to a Bézier as $min\|\mathcal{B}(t) - \mathbf{p}\|$. The closest point on the curve is denoted as $\mathcal{B}(t^{\mathbf{P}})$ (Fig. 4). We will need these definitions for writing the robot state and in the innovation step of UKF.

4) A Problematic Issue: Due to the featureless geometry of the tunnel, we cannot infer the robot position along the tunnel axis using only a range sensor and an IMU. Consider the case that the robot is flying inside a compass prohibitive, torus shaped tunnel. The robot cannot observe its global yaw angle. For this reason, we can estimate the robot yaw only with respect to the local tunnel axis. Since the robot does not maneuver aggressively, the other two rotational DOFs can be estimated using the on-board IMU. The remaining two translational DOFs are estimated with respect to the local frame.

5) The System State: The system state vector writes as

$$\mathbf{s}^\top = \left[\hat{\mathbf{g}}^\top, \mathbf{b}^\top, \mathbf{v}^\top, \omega^\top, \mathcal{S}^{0^\top}, ..., \mathcal{S}^{i^\top}, ...\right] \quad (1)$$

where $\hat{\mathbf{g}} \in \mathbb{S}^2$ is the gravity vector direction, $\mathbf{b} \in \mathbb{R}^6$ is accelerometer and gyroscope biases, $\mathbf{v}, \omega \in \mathbb{R}^3$ are the translational and rotational velocities of $\mathcal{M}$ segments all defined in $\mathcal{V}$. The size of $\mathbf{s}$ changes depending on the visible range of the tunnel by addition and removal of segments. The unconventional state vector, $\mathbf{s}$, has the minimal set of variables that fully represents the robot state and the local map in this unique environment. Finally the system state covariance is denoted as $\mathbf{\Sigma}_{ss}$.

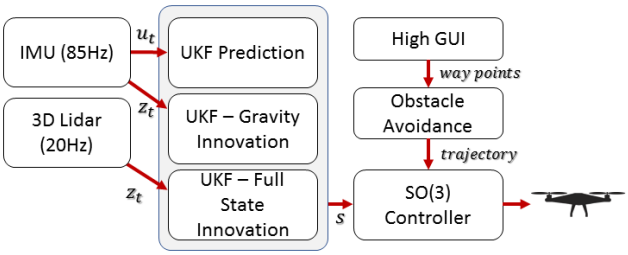


Fig. 6: System design with inputs from an IMU and a 3D lidar. A UKF integrates the information from the IMU and the lidar to obtain 5-DOF state estimate. The missing DOF is the position along the tunnel axis. The user sends high-level way point commands through a GUI. Obstacle avoidance block filters the user commands before the low level on-board SO(3) controller for safety.

6) The Robot State: The robot state is defined as $\mathbf{x}^\top = \left[\mathbf{r}^\top, \dot{\mathbf{r}}^\top, \mathbf{q}^\top\right]$ where $\mathbf{r}, \dot{\mathbf{r}} \in \mathbb{R}^2$ and $\mathbf{q} \in \mathbb{H}_1$. $\mathbf{r}$ and $\dot{\mathbf{r}}$ are defined in 2D the tunnel cross-section and $\mathbf{q}$ is defined with respect to the gravity vector and the local tunnel axis (Fig. 4).

The instantaneous reference frame with respect to which $\mathbf{r}$ and $\dot{\mathbf{r}}$ are written is $\mathcal{O}^{\mathcal{B}^i(t^{\mathcal{O}^{\mathcal{V}}})}$, $\mathcal{L}^{\mathcal{B}^i(t^{\mathcal{O}^{\mathcal{V}}})}$ (we will use $\mathcal{O}^*$ and $\mathcal{L}^*$ for clarity). These correspond to the center and the local reference frame of the tunnel at the cross-section that the robot is on. This is an intuitive choice of frame because it facilitates way point commanding and trajectory following. A good example to this is a robot commanded to fly along the centerline or at a specific offset from the tunnel center during the whole traversal. These two vectors read as

$$\mathbf{r} = -\Gamma^{\mathcal{L}^*} \mathbf{R}_{\mathcal{V}}\, \mathcal{B}^* \qquad \dot{\mathbf{r}} = -\Gamma^{\mathcal{L}^*} \mathbf{R}_{\mathcal{V}}\, \mathbf{v} \quad (2)$$

$$\Gamma = \left[\ \mathbf{0}_{2\times 1},\ \ \mathbf{I}_{2\times 2}\ \right] \quad (3)$$

The rotation matrix, $\mathbf{R}(\mathbf{q})$, corresponding to $\mathbf{q}$ is

$$\mathbf{R}(\mathbf{q}) = \left[\ \hat{\mathbf{g}} \times (\hat{\mathbf{x}}^{\mathcal{L}^*} \times \hat{\mathbf{g}}),\ \ \hat{\mathbf{x}}^{\mathcal{L}^*} \times \hat{\mathbf{g}},\ \ \hat{\mathbf{g}}\ \right]^\top \quad (4)$$

where each element is a *normalized* column vector (not shown for clarity).

IV. Methodology

A. System Overview

In this section we give a system overview in order to facilitate following the details of the proposed method. The *system state*, $\mathbf{s}$, is tracked using a constrained UKF running on the Riemannian manifold $\mathbb{R} \times \mathbb{S}^2 \times \mathbb{S}^2$. Its prediction step and measurement update are applied for every IMU and pointcloud data respectively.

The input to the *measurement model* is a chain of segments, $\{\tilde{\mathcal{S}}^i\}$, which is obtained through a non-linear constrained optimizer applied on the filtered pointcloud data, $\mathcal{P}$, and its surface normals, $\mathcal{N}$ (Fig. 4). We call $\{\tilde{\mathcal{S}}^i\}$ *measured segments*. The optimizer runs on $\mathbb{S}^2 \times \mathbb{S}^2$ to find radius, position and axis of each measured segment. This iterative process fits one segment at a time with every next initial

state being a function of the previous segment. We use Watson distribution on $\mathbb{S}^2$ (Sec. IV-F) [13] for constraining the solution space of the optimizer and avoid divergence in presence of noisy lidar data . Fig. 6 illustrates the system diagram with input/output relations.

A unique challenge of this work is the data association and innovation steps of UKF. There is no trivial distance metric between $\mathcal{M}$ and measured segments, $\{\tilde{\mathcal{S}}^i\}$. In order to overcome this problem, we use Bézier to interpolate the intermediate points in $\mathcal{M}$ with knots defined by each consecutive $\mathbf{k}^i - \hat{\mathbf{t}}^i$ pair which we denote as $\mathcal{B}^i$ (Sec. III-C.3). Then, each measured segment, $\tilde{\mathcal{S}}^i$, is compared against each spline in $\{\mathcal{B}^i\}$ for data association. Every measured segment is associated with its closest Bézier in the sense explained in Sec. III-C.3. A new spline is found by morphing the associated $\mathcal{B}^i$ so as to pass through the corresponding measured segment position and tangent with minimal distortion on the original spline (Sec. IV-B). The innovation for each segment is then calculated from the *difference* between the morphed and the original Bézier 's. These processes are illustrated in Fig. 4-5 and further details are given in the subsequent sections.

B. Morphing Cubic Bézier Splines

We use Bézier interpolation and the distance metric defined in Sec. III-C.3 for associating measured segments, $\{\tilde{\mathcal{S}}^i\}$, to map segments, $\{\mathcal{S}^j\}$. $\tilde{\mathcal{S}}^i$ is associated to the closest Bézier denoted as $\mathcal{B}^{i_*}$. This is then morphed so that it passes through $\tilde{\mathbf{k}}^i$ with its tangent parallel to $\hat{\mathbf{t}}^i$ by optimizing for a cost function. However there are infinitely many solutions to this operation. We reduce the solution to a single Bézier by applying a constraint and penalties to the optimizer. We denote the resultant morphed Bézier as $\mathcal{B}^{i_*}_m$ and its constructor segments as $\{\mathcal{S}^{i_*^-}_m - \mathcal{S}^{i_*}_m\}$. The cost function reads as $C = C_1 + C_2 + C_3$ where

$$C_1 = \left\| \mathcal{B}^{i_*}\left(t^{\tilde{\mathbf{k}}^i}\right) - \tilde{\mathbf{k}}^i \right\|^2 + \hat{\alpha}\left(1 - \left(\hat{\tilde{\mathbf{t}}}^i \cdot \hat{\mathcal{B}}^{i'_*}(t^{\tilde{\mathbf{k}}^i})\right)^2\right) \quad (5)$$

$$C_2 = \|\mathbf{k}^{i_*^-}_m - \mathbf{k}^{i_*^-}\|^2 + \|\mathbf{k}^{i_*}_m - \mathbf{k}^{i_*}\|^2 \quad (6)$$

$$C_3 = 2 - \left(\hat{\mathbf{t}}^{i_*^-}_m \cdot \hat{\mathbf{d}}\right)^2 + \left(\hat{\mathbf{t}}^{i_*}_m \cdot \hat{\mathbf{d}}\right)^2 \quad (7)$$

where $\hat{\mathbf{d}} = \mathbf{d}/\|\mathbf{d}\|$, $\mathbf{d} = \mathbf{k}^{i_*}_m - \mathbf{k}^{i_*^-}_m$. We also impose the constraint $\|\mathbf{d}\| = const$. C_2 and C_3 penalize divergence from the initial Bézier and the total curvature (equivalently curve length) respectively. This operation is illustrated in Fig. 5.

The need for defining a morphing operation emerges due to innovation step of UKF. Through the morphing operation we convert a measurement, $\tilde{\mathcal{S}}^i$, into a Bézier , $\mathcal{B}^{i_*}_m$. This is then compared against the original Bézier , $\mathcal{B}^{i_*}$, from the result of which we calculate the *innovation* in the UKF measurement update (Sec. IV-D).

C. UKF Process and Measurement Models

The filtering formulation presented here offers a solution for both local map and robot state estimation on an MAV platform flying inside a PSGC shaped tunnel environment. The platform is comprised of an IMU and a 3D lidar. IMU is used both as a control input to the process model and a measurement source for gravity vector estimation. Lidar on the other hand provides raw pointcloud data from which a parametric approximation to the tunnel is fit denoted as $\{\tilde{\mathcal{S}}^i\}$. This approximation is then used in the measurement update of the UKF.

1) Process Model: The process model $\mathcal{A}$ estimates the evolution of the system state, $\mathbf{s}$ (Sec. III-C.6), due to the control input, $\mathbf{u}$ and the process noise, $\mathbf{w}$. The form of the process models is

$$\mathbf{s}_{k+1} = \mathcal{A}(\mathbf{s}_k, \mathbf{u}_k, \mathbf{w}_k) \quad (8)$$

$$\mathbf{u}^\top = \begin{bmatrix} \mathbf{u}_a{}^\top, & \mathbf{u}_\omega{}^\top \end{bmatrix} \quad (9)$$

$$\mathbf{w}^\top = \begin{bmatrix} \mathbf{w}_a{}^\top, & \mathbf{w}_\omega{}^\top, & \mathbf{w}_{\mathbf{b}_a}{}^\top, & \mathbf{w}_{\mathbf{b}_\omega}{}^\top \end{bmatrix} \quad (10)$$

where $\mathbf{u}$ is the control input comprising of acceleration and rotational velocity data from the IMU, $\mathbf{w} \sim \mathcal{N}(0, \mathbf{Q})$ is the process noise consisting of noise in acceleration, angular velocity, acceleration bias, gyroscope bias, and $\mathbf{Q}$ is the process noise covariance.

The process model $\mathcal{A}$ is defined as

$$\hat{\mathbf{g}}_{k+1} = \mathbf{R}(\omega_k \Delta t)\hat{\mathbf{g}}_k \quad (11)$$

$$\mathbf{b}_{k+1} = \mathbf{b}_k + \mathbf{w}_{\mathbf{b},k}\Delta t \quad (12)$$

$$\mathbf{v}_{k+1} = \mathbf{v}_k - \tilde{a}_k \Delta t \quad (13)$$

$$\omega_{k+1} = \omega_k - (\mathbf{u}_\omega - \mathbf{b}_\omega + \mathbf{w}_\omega) \quad (14)$$

$$\mathbf{k}^i_{k+1} = \mathbf{R}(\omega_k \Delta t)\mathbf{k}^i_k + \mathbf{v}_k \Delta t - \frac{1}{2}\tilde{a}_k \Delta t^2 \quad (15)$$

$$\hat{\mathbf{t}}^i_{k+1} = \mathbf{R}(\omega_k \Delta t)\hat{\mathbf{t}}^i_k \quad (16)$$

where $\tilde{a} = \mathbf{u}_a - \mathbf{b}_a + \mathbf{w}_a$ and $\mathbf{R}(\omega)$ is the rotation matrix corresponding to the rotation vector ω.

2) Measurement Model: The measurement model, $\mathcal{H}$, relates a measurement value, $\mathbf{z}$, to the system state, $\mathbf{s}$. It also integrates the effect of the measurement noise, $\mathbf{v}$, into the expected measurement at a specific state.

In this work, we use two measurement models $\mathcal{H}_{\hat{\mathbf{g}}}$ and $\mathcal{H}_{\mathcal{S}}$ for gravity vector estimation using the IMU and to update the whole system state using the pointcloud data respectively. Input to the latter model is a chain measured segments, $\{\tilde{\mathcal{S}}^i\}$. $\{\tilde{\mathcal{S}}^i\}$ is obtained from $\mathcal{P}$ details of which are explained in Section IV-G.

The general form of a measurement models is

$$\mathbf{z}_k = \mathcal{H}(\mathbf{s}_k, \mathbf{v}_k) \quad (17)$$

where $\mathbf{z}$ is either $\mathbf{u}_a$ or $\tilde{\mathcal{S}}^i$. $\mathbf{v} \sim \mathcal{N}(0, \mathbf{R})$ is the sensor noise defined separately for each measurement model. The former measurement model is defined as

$$\mathcal{H}_{\hat{\mathbf{g}}} : \mathbf{u}_a = g\hat{\mathbf{g}} + \mathbf{b}_a + \mathbf{v}_{\hat{\mathbf{g}}} \quad (18)$$

where g is the gravitational acceleration and $\mathbf{v}_{\hat{\mathbf{g}}} \sim \mathcal{N}(0, \mathbf{R}_{\hat{\mathbf{g}}\hat{\mathbf{g}}})$. The latter measurement model is defined as

$$\mathcal{H}_{\mathcal{S}} : \begin{bmatrix} \mathcal{S}^{i_*^-}_m \\ \mathcal{S}^{i_*}_m \end{bmatrix} = \begin{bmatrix} \mathcal{S}^{i_*^-} \\ \mathcal{S}^{i_*} \end{bmatrix} + \mathbf{v}_{\mathcal{S}} \quad (19)$$

where the left-hand-side is the spline, $\mathcal{B}^{i_*}_m$ obtained by morphing the spline defined by the right-hand-size vector,

$\mathcal{B}^{i*}$. The details of the morphing operation is explained in Sec. IV-B. Finally, $\mathbf{v}_{\mathcal{S}} \sim \mathbf{R}_{\mathcal{SS}}$ is the measurement noise.

D. A Constrained UKF on a Riemannian Manifold

The unit vectors in the system state, $\mathbf{s}$, belong to $\mathbb{S}^2$ which prohibits use of a linear Kalman Filter (KF). For this reason we use a UKF that neither requires linear process or measurement models nor linearization of these as in the case of Extended Kalman Filter (EKF). Through application of process and measurement models on the sigma points, one can attain significantly better results compared to alternative KF variations [12]. Our UKF runs on $\mathbb{R} \times \mathbb{S}^2 \times \mathbb{S}^2$. We leave the details of UKF to the seminal paper [12] and explain issues specific to the state definition.

The state covariance, denoted as $\mathbf{\Sigma}_{ss}$, is a $12+7n$ positive square matrix where $n = 1 + N^- + N^+$. Sigma points are obtained as $\mathcal{X}_i = \mathbf{s} \pm \mathcal{W}_i$ where $\{\mathcal{W}_i\}$ are the columns of $S = \sqrt{\mathbf{\Sigma}_{ss}}$ obtained through *Cholesky* decomposition. The blocks of $\mathbf{\Sigma}_{ss}$ corresponding to $\mathbf{b}, \mathbf{v}, \omega, \mathbf{k}^i$ and ρ^i hold the corresponding uncertainties which belong to $\mathbb{R}^3$. However, the uncertainties of $\hat{\mathbf{g}}$ and $\hat{\mathbf{t}}$ are represented as 3D rotation vectors. In sigma point generation, mean and covariance calculation steps, these elements require special handling. That is, we have to define summation, subtraction and mean operation on these elements.

We define the summation of a unit vector, $\hat{\mathbf{a}}$, with a rotation vector, $\mathbf{b}$, as $\hat{\mathbf{a}} + \mathbf{b} = \mathbf{R}(\mathbf{b})\hat{\mathbf{a}}$ where $\mathbf{R}(\mathbf{b})$ is the rotation matrix corresponding the quaternion $\mathbf{q} = \left[\ \cos\left(\|\mathbf{b}\|/2\right),\quad \sin\left(\|\mathbf{b}\|/2\right)\ {}^{\mathbf{b}}/_{\|\mathbf{b}\|}\ \right]$. Similarly the subtraction operation is defined as $\hat{\mathbf{a}} - \mathbf{b} = \mathbf{R}(-\mathbf{b})\hat{\mathbf{a}}$. The mean of a set of unit vectors, $\{\hat{\mathbf{a}}\}$, is $\hat{\bar{\mathbf{a}}}$ which minimizes $\sum_{\{\hat{\mathbf{a}}\}} \hat{\mathbf{a}} \times \bar{\mathbf{a}}$ [21].

The operations defined above are also used for obtaining the innovation term of UKF. Each measured segments is converted into a Bézier through morphing (Sec. IV-B). The constructor segments of the morphed spline, $\mathcal{S}_m^{i_*^-}$ and $\mathcal{S}_m^{i_*}$, is *subtracted* in order from their corresponding original segments, $\mathcal{S}^{i_*^-}$ and $\mathcal{S}^{i_*}$, to obtain the innovation on the two segments.

The innovation of UKF does not respect the constant distance constraint, $\|\mathcal{S}^i - \mathcal{S}^{i^-}\| = \ell^i$, between consecutive segments (Sec. III-C.2). We use the method explained Section 3 of [22] to integrate the effect of this non-linear constraint on the segment positions at the end of each UKF cycle.

E. Surface Normal Estimation

Surface normals play an important role in the proposed method since the performance of the measurement model is based on accurate normal estimation. We use the method presented in [23] for surface normal estimation. This method uses Kd-Tree's for neighborhood search and analyze the spread of the point in the 3D space. The eigenvector corresponding to the smallest eigenvalue is taken as the surface normal represented by the neighboring points. The ratio of this to the summation of the three eigenvalues is taken as the curvature, κ, at the point of interest [15]. We apply a voxel filter on $\mathcal{P}$ to reduce the point count to save CPU time. We use Watson distribution on $\mathbb{S}^2$ for analyzing surface normals and eliminate outliers as explain in the next section.

F. Watson Distribution on $\mathbb{S}^2$

The performance of the measurement model (Sec. IV-G) heavily depends on the accuracy of surface normals. Because of this we apply outlier elimination on the raw normal data using directional statistical analysis tools. Specifically, we use Watson distribution (Section 9.4 in [13]) to estimate the null space of the surface normals in presence of noisy data and inaccurate normal estimation. Density of Watson distribution is written as

$$\mathcal{W}(\pm\mathbf{x}, \mu, \kappa) = M\left(\frac{1}{2}, \frac{p}{2}, \kappa\right)^{-1} exp\left(\kappa(\mu^\top \mathbf{x})^2\right) \tag{20}$$

where μ is the mean direction, κ is the concentration value and $M\left({}^1/_2, {}^p/_2, \kappa\right)$ is the Kummer function which is

$$M\left(\frac{1}{2}, \frac{p}{2}, \kappa\right) = B\left(\frac{p-1}{2}, \frac{1}{2}\right)^{-1} \int_{-1}^{1} e^{\kappa t^2} (1-t^2)^{(p\,-\,3)/2} dx \tag{21}$$

with $p = 3$ for the $\mathbb{S}^2$ case. Watson takes the form of a bipolar distribution for $\kappa > 0$ and becomes a girdle distribution for $\kappa < 0$. Different synthetic data and the results of outlier elimination are shown in Fig. 7.

Either due to inaccurate surface normals or violation of the PSGC environment assumption, often the normals inside a segment do not point towards the tunnel axis which in turn disturb the optimal solution of the segment measurement (Sec. IV-G). In order to clean the data from spurious normals, we find the most likely *girdle* distribution by iteratively eliminating less likely data points. We use the techniques explained in Section 10.3.2 of [13] to infer the model parameters (μ, κ). Using *F-test* as explained in Section 10.7.5 of [13] we also get a confidence cone around μ (Fig. 7). Confidence cone is basically a function of the eigenvalues of the scatter matrix of the inlier surface normals [13]. Normals outside the 99% confidence region of the cumulative Kummer function are eliminated. This process is iterated until convergence. Results of this process are used for initializing and constraining the non-linear optimizer explained in Sec. IV-G as well as eliminating outlier from $\mathcal{P}$ and $\mathcal{N}$.

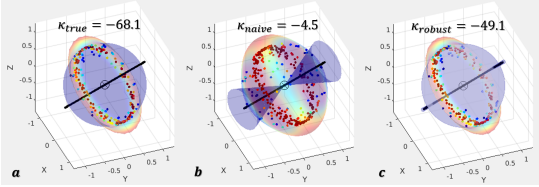
The Watson distribution fit for $\mathcal{S}^i$ is denoted by $\mathcal{W}^i$ and the corresponding mean and concentration parameters are denoted as μ^i, κ^i. Finally the apex angle of the confidence region is called θ^i.

G. Non-Linear Optimization for Measured Segments

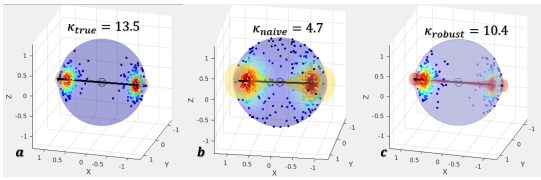
We formulate the segment measurement as a constrained optimization problem running on $\mathbb{S}^2 \times \mathbb{S}^2$. This seeks for the optimal solution for the cost function $C = C_1 + C_2$ where

$$C_1 = \frac{1}{|\mathcal{P}^i|} \sum_{\mathbf{p} \in \mathcal{P}^i} |\bar{\mathbf{r}} - \mathbf{r}_n| \tag{22}$$

$$C_2 = \frac{1}{|\mathcal{N}^i|} \sum_{\hat{\mathbf{n}} \in \mathcal{N}^i} \left|\hat{\mathbf{n}} \cdot \hat{\mathbf{x}}^i\right| \tag{23}$$



(a) Girdle distribution, $\kappa < 0$



(b) Bipolar distribution, $\kappa > 0$

Fig. 7: Synthetic data on $\mathbb{S}^2$ sampled from a girdle and a bipolar distribution. The middle and right plots show the naïve and robust Watson fitting results with 30% uniform noise added. Confidence region of the mean direction is shown with a cone.

where

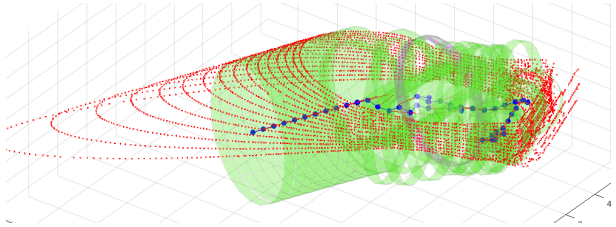
$$\mathbf{r}_n = (\mathcal{I} - \hat{\mathbf{x}}^i \, \hat{\mathbf{x}}^{i^\top})(\mathbf{p} - (\mathbf{k}^{i^-} + \tilde{\ell}^i \hat{\alpha}^i)) \tag{24}$$

$$\bar{\mathbf{r}} = \frac{1}{|\mathcal{P}^i|} \sum_{\mathbf{p} \in \mathcal{P}^i} \mathbf{r}_n, \tag{25}$$

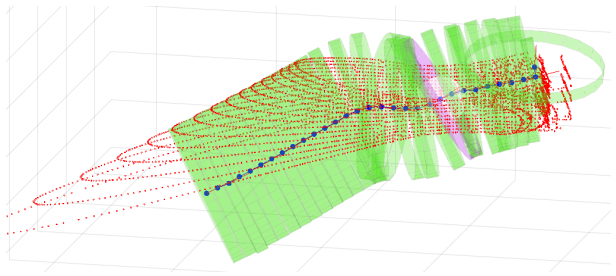
are respectively the distance of $\mathbf{p}$ from the current local axis estimate of $\tilde{\mathcal{S}}^i$, and $\bar{\mathbf{r}}$ is the mean radius. $\mathcal{P}^i$ and $\mathcal{N}^i$ the subset of the original pointcloud data within a bounding box centered at $\mathcal{O}^i$ and aligned with the local frame $\mathcal{L}^i$. These are filtered to exclude high curvature points which are very likely to be outlier such as scaffolding or operators (Sec. IV-E). $\tilde{\ell}^i$ is the fixed chord length between $\tilde{\mathcal{S}}^{i^-}$ and $\tilde{\mathcal{S}}^i$, and $\hat{\alpha}^i$ is the chord direction. This formulation allows for keeping the chord length, $\tilde{\ell}$, constant at the cost of optimizing for the current segment position on $\mathbb{S}^2$ rather than $\mathbb{R}^3$. Furthermore, by imposing an angle constraint on the orientation of $\hat{\alpha}^i$, we easily avoid divergence of the optimizer.

The free variables of the optimization are $\hat{\alpha}$ and $\hat{\mathbf{x}}$. The initial values are inherited from the previous segment such that $\hat{\alpha} = \hat{\mathbf{x}}^{i^-}$ and $\hat{\mathbf{x}}^i = \hat{\mathbf{x}}^{i^-}$. The constraints on these values are derived from the Watson distribution fit of the previous segment, $\mathcal{W}^{i^-}$. Since the tunnel is a smooth structure, we can assume that the maximum amount of bending between adjacent segments is a function of the confidence region of $\mathcal{W}^{i^-}$. Hence we impose the following constraints $\hat{\alpha}^i \cdot \mu^{i^-} \geq \cos(c_{\hat{\alpha}} \, \theta^{i^-})$ and $\hat{\mathbf{x}}^i \cdot \mu^{i^-} \geq \cos(c_{\hat{\mathbf{x}}} \, \theta^{i^-})$ for $c_{\hat{\alpha}} \geq c_{\hat{\mathbf{x}}} > 0$.

The optimizer uses the Marquardt-Levenberg method with adaptive step size to avoid saddle points or non-positive Hessian matrices equivalently. The update is written as $\left[\, \Delta\hat{\alpha}^\top, \quad \Delta\hat{\mathbf{x}}^\top \, \right]^\top = -\lambda \; \mathbf{H}(C)^{-1} \; \nabla g(C)$ where λ is the step size, $\mathbf{H}$ and ∇g are the Hessian and the gradient operators. Each update is first projected onto $\mathbb{S}^2$ at the most



(a) Without filtering



(b) With filtering

Fig. 8: The robot is flying close to the gate in a penstock at Center Hill Dam, TN where PSGC assumption does not hold. The segment measurements fail since the surface normals due to the gate pollute $\mathcal{N}$. Two figures show the estimator performance with and without filtering based on Watson analysis.

recent position as

$$\hat{\alpha}_{new} = \hat{\alpha} + (\mathcal{I} - \hat{\alpha}\hat{\alpha}^\top)\Delta\hat{\alpha} \tag{26}$$

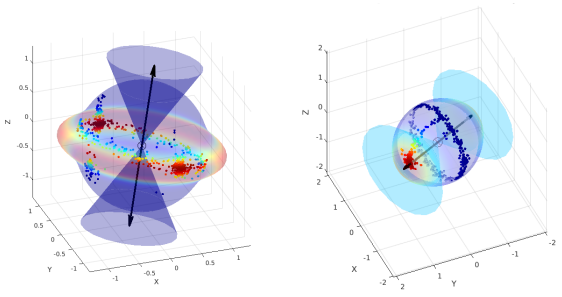
$$\hat{\mathbf{x}}_{new} = \hat{\mathbf{x}} + (\mathcal{I} - \hat{\mathbf{x}}\hat{\mathbf{x}}^\top)\Delta\hat{\mathbf{x}} \tag{27}$$

followed by normalization (not shown). The radius is automatically estimated as in Equ. 25 once the optimization converges. The uncertainty of the segment is taken as the inverse of the Hessian. Finally, after convergence, a Watson is fit for the current segment to define the constraint on the next segment.

When the robot flies close to a non-PSGC region, the optimizer may diverge due to the failure of the circularity assumption. In order to avoid this, we add samples to $\mathcal{N}^i$ drawn from a girdle distribution with $\kappa \ll 0$ and mean $\tilde{\hat{\mathbf{t}}}^i$. This biases $\mathcal{W}^i$ towards a proper girdle distribution as a result of which the *next* segment is initialized at a better point with better optimization constraints. This procedure has no affect if the PSGC assumption holds. However, especially when the robot is close to gate, a significant percentage of $\mathcal{N}^i$ is due to the surface normals at the gate hence the resultant distribution tends to be bipolar. We have seen that the optimizer performs much better in such situations. Fig. 8 shows a comparison of these two cases. Also Fig. 3 shows successful segment measurements without filtering at a region where PSGC holds.

V. Experiments

We would like to challenge our on-board estimator when the PSGC environment assumption does not hold. Thus we present results for the dataset collected in one of the penstocks at Center Hill Dam, TN. There are three main



(a) Without filtering (b) With filtering

Fig. 9: These two plots show the two Watson distributions corresponding to the two cases shown in Fig. 8 when the robot is flying very close to the gate at Center Hill Dam, TN. When filtering and biasing is not activated, the error accumulates and fails the estimator. However, when the we effectively use the filter on $\mathbb{S}^2$, outliers are eliminated and tunnel axis (shown with line segments in Fig. 8) is estimated with greater accuracy. The arrows in these plots correspond to the tunnel axis.

cases that we test in this work, which are the estimator performance during flight along the horizontal and inclined sections. Lastly, we test performance when the robot is very close to the gate. The results we obtained for the latter case is of significant importance since the PSGC assumption is heavily violated due to the geometry change in the vicinity of the gate. The tunnel takes a rectangular prism shape and the 3D lidar only sees the side walls and the gate hence the surface normals focus to a bipolar distribution which does not offer cylindrical shape at all. For all cases, we highlight the performance of the estimator in presence of noise and violations on the PSGC environment assumption. A comparison between two runs with and without smoothness constraint is presented in Fig. 8 to highlight the performance difference. Fig. 9 shows the surface normals and the corresponding axis estimates for the two cases.

Since we do not have ground truth data, we evaluate the results visually. We refer to the alignment of segments with the raw pointcloud data in the visualization tool.

VI. Conclusion

This work presents a new filter designed for autonomous flights of MAVs inside piecewise-smooth-generalized-cylinder shaped tunnel-like environments. We show results for the performance on a dataset collected in a penstock at Center Hill Dam, TN. The penstock is modeled as a deformable chain of generalized cylindrical segments which allows for imposing smoothness constraint on the map. Using the spherical Watson distribution $\mathbb{S}^2$, we do outlier elimination to robustify the estimator. A UKF running on a manifold integrates information from a 3D lidar and an IMU.

References

[1] N. Michael, D. Mellinger, Q. Lindsey, and V. Kumar, "The GRASP multiple micro-UAV testbed," *IEEE Robot. Autom. Mag.*, vol. 17, no. 3, pp. 56–65, sep 2010.

[2] S. Weiss, D. Scaramuzza, and R. Siegwart, "Monocular-SLAM-based navigation for autonomous micro helicopters in GPS-denied environments," *J. F. Robot.*, vol. 28, no. 6, pp. 854–874, nov 2011.

[3] D. Cole and P. Newman, "Using laser range data for 3D SLAM in outdoor environments," in *Proc. 2006 IEEE Int. Conf. Robot. Autom. 2006. ICRA 2006.* IEEE, 2006, pp. 1556–1563.

[4] G. Loianno, C. Brunner, G. McGrath, and V. Kumar, "Estimation, Control, and Planning for Aggressive Flight With a Small Quadrotor With a Single Camera and IMU," *IEEE Robot. Autom. Lett.*, vol. 2, no. 2, pp. 404–411, apr 2017.

[5] S. Liu, M. Watterson, S. Tang, and V. Kumar, "High speed navigation for quadrotors with limited onboard sensing," in *Proc. - IEEE Int. Conf. Robot. Autom.*, vol. 2016-June. IEEE, may 2016, pp. 1484–1491.

[6] T. Tomic, K. Schmid, P. Lutz, A. Domel, M. Kassecker, E. Mair, I. Grixa, F. Ruess, M. Suppa, and D. Burschka, "Toward a fully autonomous UAV: Research platform for indoor and outdoor urban search and rescue," *IEEE Robot. Autom. Mag.*, vol. 19, no. 3, pp. 46–56, sep 2012.

[7] T. Özaslan, K. Mohta, J. Keller, Y. Mulgaonkar, C. J. Taylor, V. Kumar, J. M. Wozencraft, and T. Hood, "Towards fully autonomous visual inspection of dark featureless dam penstocks using MAVs," in *2016 IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, oct 2016, pp. 4998–5005.

[8] E. Eade and T. Drummond, "Scalable Monocular SLAM," in *CVPR*, vol. 1. IEEE, 2006, pp. 469–476.

[9] J. Zhang and S. Singh, "Visual-lidar odometry and mapping: Low-drift, robust, and fast," in *Robot. Autom. (ICRA), 2015 ...*. IEEE, may 2015, pp. 2174–2181.

[10] S. Shaojie, "Autonomous Navigation in Complex Indoor and Outdoor Environments with Micro Aerial Vehicles," Ph.D. dissertation, 2014.

[11] S. M. LaValle, *Planning Algorithms*. Cambridge, U.K.: Cambridge University Press, 2006.

[12] R. Van der Merwe and E. Wan, "The square-root unscented Kalman filter for state and parameter-estimation," in *2001 IEEE Int. Conf. Acoust. Speech, Signal Process. Proc. (Cat. No.01CH37221)*, vol. 6. IEEE, pp. 3461–3464.

[13] K. V. Mardia and P. E. Jupp, *Directional Statistics [Hardcover]*. J. Wiley, 1999.

[14] T. Özaslan, S. Shen, Y. Mulgaonkar, N. Michael, and V. Kumar, "Inspection of Penstocks and Featureless Tunnel-like Environments Using Micro UAVs," in *F. Serv. Robot. 2013*, ser. Springer Tracts in Advanced Robotics, L. Mejias, P. Corke, and J. Roberts, Eds., vol. 105. Brisbane: Springer International Publishing, jan 2015, pp. 123–136.

[15] T. Özaslan, G. Loianno, J. Keller, C. J. Taylor, V. Kumar, J. M. Wozencraft, and T. Hood, "Autonomous Navigation and Mapping for Inspection of Penstocks and Tunnels with MAVs," *IEEE Robot. Autom. Lett.*, vol. 2, no. 3, pp. 1740–1747, jul 2017.

[16] A. Ortiz, F. Bonnin-Pascual, and E. Garcia-Fidalgo, "Vessel Inspection: A Micro-Aerial Vehicle-based Approach," *J. Intell. Robot. Syst.*, vol. 76, no. 1, pp. 151–167, sep 2014.

[17] M. Burri, J. Nikolic, C. Hürzeler, G. Caprari, C. Hurzeler, G. Caprari, R. Siegwart, C. Hürzeler, G. Caprari, C. Hurzeler, G. Caprari, and R. Siegwart, "Aerial service robots for visual inspection of thermal power plant boiler systems," *Appl. Robot. Power Ind. (CARPI), 2012 2nd Int. Conf.*, pp. 70–75, 2012.

[18] P. Gohl, M. Burri, S. Omari, J. Rehder, J. Nikolic, M. Achtelik, and R. Siegwart, "Towards autonomous mine inspection," in *Proc. 3rd Int. Conf. Appl. Robot. Power Ind. CARPI 2014*, 2015.

[19] J. Zhang and S. Singh, "Low-drift and real-time lidar odometry and mapping," *Auton. Robots*, vol. 41, no. 2, pp. 401–416, feb 2017.

[20] A. S. Huang, A. Bachrach, P. Henry, M. Krainin, D. Maturana, D. Fox, and N. Roy, "Visual Odometry and Mapping for Autonomous Flight Using an RGB-D Camera," in *Proc. Int. Symp. Robot. Res.*, 2011, vol. 100, pp. 235–252.

[21] E. Kraft, "A quaternion-based unscented Kalman filter for orientation tracking," in *Proc. 6th Int. Conf. Inf. Fusion, FUSION 2003*, vol. 1, 2003, pp. 47–54.

[22] D. Simon, "Kalman filtering with state constraints: a survey of linear and nonlinear algorithms," *IET Control Theory Appl.*, vol. 4, no. 8, pp. 1303–1318, aug 2010.

[23] R. B. Rusu, "Semantic 3D Object Maps for Everyday Manipulation in Human Living Environments," Ph.D. dissertation, TECHNISCHE UNIVERSITÄT MÜNCHEN, nov 2010.

2017 Final Report Metrics

Task Number: BPP-SUPP-16-17-4-CNC

Metrics

Notes: *Please list metrics for the reporting quarter only and only Publications that are accepted and searchable on web.*

Awards and/or Recognitions this Quarter

-

Refereed/Journal Papers accepted this quarter

T. Ozaslan, G. Loianno, J. Keller, CJ. Taylor, V. Kumar, JM. Wozencraft and T. Hood, "Robust State Estimation for Tunnel Inspection with MAVs" in IEEE Robotics and Automation Letters (RA-L), 2017

Conference Presentations/proceedings presented this quarter

T. Ozaslan, G. Loianno, J. Keller, CJ. Taylor, V. Kumar, "Spatio-Temporally Smooth Local Mapping and State Estimation Inside Generalized-Cylindrical Tunnels with Inspection MAVs" in IEEE International Conference on Robotics and Automation (ICRA), 2018 *(under review)*

Joint Papers with ARL presented or accepted this quarter – duplication with journals/ conferences expected.

Invention Disclosures for this quarter

Patent Applications Filed this quarter

-Patents Awarded this quarter

PhDs Award this quarter:

Masters Awarded this quarter

Students supported fully or partially (by name)

MS: *Qiaoyu Zhang*
PhD: Tolga Ozaslan, James Keller, Spyridon Leonardos, Yash Mulgaonkar,
Post Docs: Jnaneshwar Das

MAST CTA Member Rotations into ARL

ARL Researchers Working at Member Sites

Workshops and Seminars conducted:

MAST Quarterly Metrics Summary Table (Numbers in this table should correspond to the list above):

MAST Final REPORT METRICS TABLE (The data entered into this table should be for the reporting quarter only)	
Awards (Miscellaneous Recognitions and Awards)	**0**
Refereed/Journal papers	**1**
Conference presentations/proceedings	**1**
Joint Papers with ARL	**0**
Invention disclosures	**0**
Patent applications filed	**0**
Patents awarded	**0**
PhDs Award	**0**
Masters Awarded	**0**
Post Docs Supported	**1**

Master Degree Students Supported	1
Undergraduate Students Supported	0
CTA Member rotations into ARL	0
Visiting Researchers from ARL to the Member Sites	0
Employees hired by ARL from CTA students supported	0

MAST Collaborative Technology Alliance – Capstone Publication Metrics
November 1, 2015 through April 2017

Aggressive Flight using Approximate Inverse Depth Models

Charles Richter and Nicholas Roy
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA 02139
Email: {car, nickroy}@mit.edu

Shreyansh Daftry and Larry Matthies
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109
{Shreyansh.Daftry, lhm}@jpl.nasa.gov

***Abstract*—The goal of this project was to enable high speed obstacle avoidance in cluttered environments without requiring full path planning through a detailed geometric world model. In previous years, we extended our methods for learning to predict the utility of future sensor measurements, enabling an autonomous robot to plan paths that enable higher average speed by seeing better into unknown regions of the environment. This used 2D occupancy grid world models created from single-axis scanning lidar data. We then developed a new method to learn to predict collision probabilities from monocular camera images directly. In the final year, we developed a perception pipeline to augment the monocular imagery with stereo vision-based depth maps and visual-inertial state estimation at high frame rates; this enabled the use of depth perception, learning, and collision checking/prediction for dynamically feasible trajectories with image-space representations, without the need for volumetric world models.**

I. Introduction

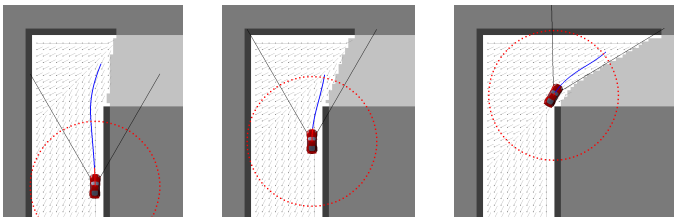
Robot obstacle avoidance in cluttered, unknown environments is typically done with motion planning algorithms that use deliberative search for collision-free trajectories in 3D representations of the environment, built onboard the robot from depth sensor data. Many variants exist, but a generate-and-test paradigm is common, where candidate trajectories are tested for potential collisions using discretized, volumetric representations of the world [6]. For small robots with limited computational power, which is the scope of the MAST program, this paradigm presents challenges with the size, weight, power, and computational throughput of the requisite sensing and processing architecture, especially as one attempts to scale up to high velocity motion. Furthermore, existing paradigms have a tendency to produce conservative motion plans, where desires for guaranteed safe behavior lead to trajectories that limit the ability of sensors to see into unknown space, which in turn limits the speed of safe travel.

We addressed these challenges in stages. The first stage retained traditional 3D world models, but explored learning-based methods to generate trajectories that give a better view into unknown areas ahead of the robot, which led to shorter overall traverse time. To facilitate safe and efficient experimentation with high speed motion, we conducted experiments using a small RC robot car equipped with a Hokuyo lidar, using an occupancy grid world model. Our longer term goals included attempting to reduce computational cost by migrating away from the use of fully 3D world models to the use of images and image-based depth representations, such as those pioneered by [7], [8]. This is a major paradigm shift, which we explored in two steps. First, we adapted and tested our learning-based approach to trajectory generation to use input from just images, where neural nets are trained to generate trajectories. Experiments showed promise for this approach. That step did not include depth information; to enable that transition, we developed a high-speed perception pipeline that includes frame rate depth perception, image-based temporal fusion, and visual-inertial motion estimation. Portions of this pipeline have been tested in our prior work that does not use learning as part of trajectory generation [9], [10]. Finally, we have taken steps toward extending and integrating this pipeline with learning-based trajectory generation.

Section II discusses learning and navigation using an occupancy grid built with lidar data. Section III describes our first step away from using 3D world models in real-time by developing methods to train a neural net to predict collisions directly from low resolution, monocular images. Lidar or an RGBD depth camera is still used to assist training and to estimate vehicle motion. Section IV describes how we leveraged and extended work from MAST task SPP-16-17-1.3 to provide a real-time depth perception pipeline with image space depth representation and visual-inertial motion estimation as a potential successor to the use of the Hokuyo lidar and occupancy grids. Section V describes how we have combined depth perception and appearance in an image-based approach that uses near-to-far learning to prove trajectories planned a long range.

II. Learning to Predict the Utility of Future Measurements

Robot navigation in unknown environments is often performed with a receding-horizon approach, where a motion planner selects a trajectory of some length that makes progress toward the goal while avoiding observed obstacles. Progress is often measured by a heuristic estimate of the remaining cost-to-goal from the end of the planned action. This approach can be represented mathematically with the following optimization



(a) 4 m/s. (b) 1 m/s. (c) 1 m/s.

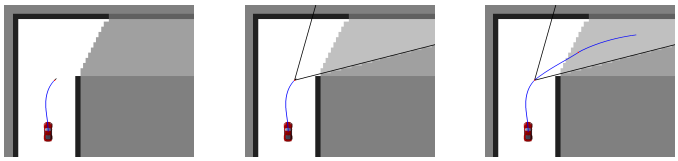
Fig. 1: Navigating with a simple shortest-path heuristic, the robot greedily keeps to the inside of the turn since it is nearer to the goal. This position occludes the free space around the corner, forcing the robot to slow down to maintain a safe stopping distance. Chosen actions a_t are drawn in blue, emergency-stop actions (lengths proportional to the square of speed) are green, and both are constrained by $g(b_t, a_t)$ to lie within known free space. Black lines indicate the field-of-view.

problem, which is repeatedly re-solved online to select the optimal action a_t^* as the robot moves through the environment:

$$a_t^* = \underset{a_t \in A}{\operatorname{argmin}} \; j(b_t, a_t) + h(b_t, a_t), \;\; \text{s.t.} \;\; g(b_t, a_t) = 0. \quad (1)$$

Here, a_t is an action to be executed at time t, chosen from a set A of possible actions. The robot's belief b_t contains its knowledge of its own configuration at time t and a partial map built from sensor measurements up to time t. The total estimated cost of choosing action a_t, given b_t, is the sum of the action cost $j(b_t, a_t)$ and the heuristic estimate $h(b_t, a_t)$ of the cost remaining beyond the end of a_t. In this work, we focus on the minimum-time navigation problem, so $j(b_t, a_t)$ returns the time duration of a_t and $h(b_t, a_t)$ estimates the remaining time to reach the goal after action a_t has been completed. Finally, to guarantee safety, we use a collision constraint $g(b_t, a_t)$, which returns 1 if action a_t intersects either obstacles or unknown regions in the belief, or leads to a state for which collision or entering unknown regions would be inevitable [1], and returns 0 otherwise. We enforce this constraint by ensuring the existence of an "emergency-stop" trajectory that could bring the robot to a stop from the end of action a_t without intersecting an obstacle or unknown region of the map. This collision constraint creates a crucial interaction between the feasibility of actions and the parts of the map that have been observed thus far. The robot is constrained to select trajectories that lie entirely within the known free space of the map, at speeds that ensure it can stop before hitting an obstacle or entering the unknown.

In an unknown map, it is very difficult or impossible to accurately estimate the remaining cost-to-goal—since doing so would amount to solving a difficult POMDP—so the heuristic $h(b_t, a_t)$ is inherently a simplified approximation. Common approximations are often based on following the apparent shortest path to the goal at a fixed speed. In this work, we define a shortest-path cost-to-goal estimator $d(b_t, a_t)$, which computes the time needed to reach the goal from the end of a_t, respecting the obstacles in belief b_t, assuming holonomic kinematics, and assuming that speed is held constant at the initial speed of a_t. We implement $d(b_t, a_t)$ using Dijkstra's



(a) (b) (c)

Fig. 2: (a)-(c) Training sequence: (a) Random robot configuration is sampled, with belief b_0 and feasible candidate action a_0; (b) the sensor is simulated from the end of a_0, yielding belief b_1 containing additional observed free space; (c) next action a_1 is selected to minimize cost, given the newly observed free space in belief b_1. The choice of a_1 will be used to compute the label of this data point.

algorithm on a 2D graph of nodes connected to their 16 nearest neighbors. Using this shortest-path function as the heuristic amounts to setting $h(b_t, a_t) = d(b_t, a_t)$ in equation (1).

To demonstrate the problems with this approach, Figure 1 illustrates a sequence of trajectories planned by a simulated car approaching a blind corner using the shortest-path heuristic. This heuristic greedily guides the robot toward the inside of the corner, where it is unable to see the free space ahead. Without visibility around the corner, the local planner is forced to select an action that slows down from 4 m/s to 1 m/s to preserve sufficient stopping distance.

This common heuristic ignores the fact that the constraints imposed by the current unknown regions of the map may actually be lifted when future observations are taken and more free space is added to the map. A more accurate heuristic might have guided the robot along a slightly longer path, sacrificing some immediate progress in exchange for improved future visibility into the unknown regions of the map. Greater visibility might, in turn, enable faster actions in future planning, reducing overall time-to-goal. Our objective in the next section will be to train a function from data that can improve the accuracy of our heuristic in exactly this way.

A. *Problem and Technical Approach*

The heuristic $h(b_t, a_t)$ approximates the expected remaining cost-to-goal after taking action a_t, given belief b_t. However, the true optimal value of this quantity represents the solution to an intractable POMDP. Rather than employing POMDP solution techniques, our approach is instead to use machine learning to somehow approximate $h(b_t, a_t)$ from training examples. While we could consider attempting to learn the entire global heuristic from data, this is not likely to be a predictable quantity due to the wide range of possible map geometries and sizes. Instead, we observe that $d(b_t, a_t)$ gives reasonable guidance much of the time, but is missing specific information about the effects of possible map observations in the immediate future, which is more local and hence predictable than the entire heuristic. Therefore, we will model the effects of possible future observations in the form of a correction to the shortest-path heuristic. Let $h^*(b_t, a_t)$ represent the expected

cost-to-goal under the optimal policy[1]. Our learned function $f_h(b_t, a_t)$ is intended to model the difference between the shortest path heuristic and this true optimal cost-to-goal:

$$f_h(b_t, a_t) \approx h^*(b_t, a_t) - d(b_t, a_t). \tag{2}$$

By capturing this effect from training data, we avoid the need to explicitly model the environment distribution and search over the vast space of observations. And, while we cannot compute $h^*(b_t, a_t)$, even offline during training, we can locally approximate it using a one-step look ahead technique, which we describe next.

1) Training: We train the model $f_h(b_t, a_t)$ using a dataset of labeled belief-action pairs, which are intended to represent realistic scenarios the robot could encounter. We generate each data point in simulation by first randomly sampling a robot configuration within a training map. We then generate a belief b_0 from the point of view of the sampled robot configuration. Having sampled a configuration and generated the initial belief, we then randomly select a feasible action a_0 the robot could execute, subject to the constraint $g(b_0, a_0) = 0$, which states that the action is collision-free and allows the robot room to stop within known free space (Figure 2a). Next, we simulate the sensor, with the actual field-of-view, from the *end* of a_0 and incorporate the measurement into b_0 to form a new updated belief, b_1 (Figure 2b). Then, we select the next action, a_1, from the end of a_0, that makes the most progress toward the goal given the updated belief b_1 and subject to the constraint $g(b_1, a_1) = 0$ (Figure 2c). Progress is measured with a calculation of $d(b_1, a_1)$, from the end of a_1. The training label is then computed using the optimal choice for a_1:

$$y(b_0, a_0) = \min_{a_1 \in A} [j(b_1, a_1) + d(b_1, a_1)] - d(b_0, a_0). \tag{3}$$

Thus, while we cannot compute $h^*(b_t, a_t)$ even offline during training, our training labels still approximate the desired result of capturing the deviation from the shortest-path heuristic that is due to potential future observations, and subsequent actions the planner could take, contingent upon those observations.

Finally, we compute a low-dimensional set of features, $\phi(b_t, a_t)$ for each belief-action pair that capture the useful predictive information. For this work, we use two features based on the boundary between known free space and the unknown, which we refer to as the "frontier": (1) the fraction of the map frontier in b_t that will be visible from the robot configurations along action a_t, and (2) the average distance from states along the action a_t to all frontiers in the map.

2) Prediction and Planning with the Learned Model: Our dataset of N data points has the form $D = \{(\phi_1, y_1), \ldots, (\phi_N, y_N)\}$ where indices in this case represent different data points (not time t) and we write ϕ_i as shorthand for $\phi(b, a)$ for the ith data point. We use the Nadaraya-Watson estimator to make predictions at runtime:

$$f_h(b_t, a_t) = \frac{\sum_{i=1}^{N} K(\phi(b_t, a_t) - \phi_i) y_i}{\sum_{i=1}^{N} K(\phi(b_t, a_t) - \phi_i)}. \tag{4}$$

Our proposed heuristic function is then the sum of the shortest-distance heuristic and the learned function: $h(b_t, a_t) = d(b_t, a_t) + f_h(b_t, a_t)$. Putting it all together, our planner repeatedly re-solves the following optimization problem:

$$a_t^* = \underset{a_t \in A}{\text{argmin}} \; j(b_t, a_t) + d(b_t, a_t) + f_h(b_t, a_t), \;\; \text{s.t.} \;\; g(b_t, a_t) = 0. \tag{5}$$

To measure the performance of our approach, we will compare it to a "baseline" planner, which uses a simpler control law that is identical to (5), except that it omits $f_h(b_t, a_t)$ and instead uses the simpler heuristic $h(b_t, a_t) = d(b_t, a_t)$.

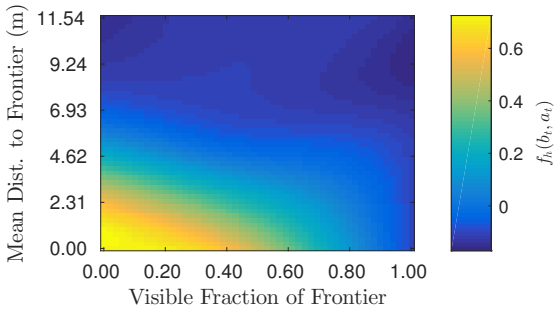


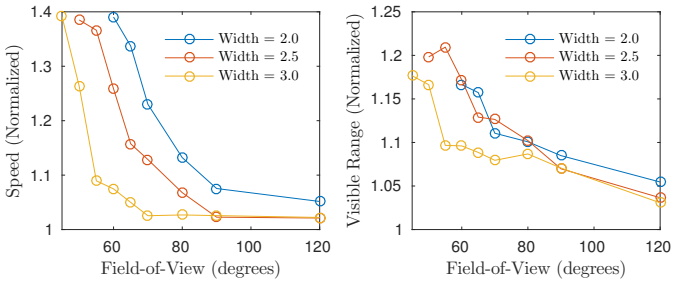
Fig. 3: Learned model $f_h(b_t, a_t)$ for 120° field-of-view in 3 m wide hallway.

B. Results for Measurement Utility Prediction

To evaluate our method, we trained a separate model for each combination of sensor field-of-view (from 45° to 120°) and hallway width (2 m, 2.5 m and 3 m). For each of these combinations, we conducted simulation trials in 50 randomly-generated hallway environments [5], recording the performance of our planner as well as that of a "baseline" planner, described above. We selected the hallway environment type because hallways are very common and also challenging for high-speed navigation due to the visual occlusion at turns. We selected the range of sensor fields-of-view to be relevant for widely used RGBD sensors such as the Microsoft Kinect and Asus Xtion, which produce point clouds of 57° and 58° horizontal field-of-view, respectively, as well as standard and wide-angle cameras.

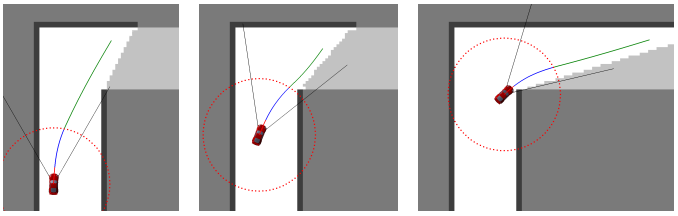
1) Learned Model: Figure 3 shows an example of a learned model for a 120° field-of-view sensor in a 3 m wide hallway. Intuitively, the highest cost belief-action pairs occur when actions are located very near to the map frontier while simultaneously offering low frontier visibility (lower left corners of plots in Figure 3). This scenario occurs when the robot approaches a corner hugging the inside wall of the hallway, which enables it to get close to the frontier while keeping the frontier mostly occluded by the wall. In fact, this costly behavior is characteristic of the baseline planner.

[1]Following the POMDP problem formulation developed in [3], this true optimal expected cost would be defined as $h^*(b_t, a_t) = \sum_{s_{t+1}} P(s_{t+1}|b_t, a_t) V^*(s_{t+1})$. We omit a detailed discussion of the POMDP formulation of this problem for brevity, but note that it provides the basis for our mathematical approach and approximations.



(a) (b)

Fig. 4: Simulation results: (a) average speed, and (b) average maximum visible range. Averaged results are normalized by the corresponding performance of the baseline planner in the same environments. Plot color indicates hallway width.



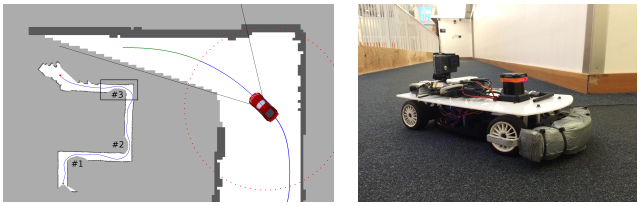
(a) 6 m/s. (b) 4 m/s. (c) 6 m/s.

Fig. 5: Example planning sequence using our method, which guides the robot wide around the corner, gaining visibility and maintaining a higher speed than the baseline planner depicted in Figure 1.

Our model predicts a reduction in cost if we select actions that either increase distance from the frontier or visibility of the frontier, or both. When the robot is approaching a blind corner, these two features both suggest approaching the corner for a wide turn, not only to increase available stopping distance between the robot position and the frontier, but also to take an observation of more occluded space, thereby pushing the frontier farther ahead.

2) Simulation Speed as a Function of Visibility: Figure 4a illustrates the average simulated navigation speeds of our method for various fields-of-view and hallway widths, normalized by the performance of the baseline planner in the same environment. In each hallway width, our solution is up to 1.4 times faster than the baseline planner. Figure 4b quantifies the greater environment visibility achieved by our method. We recorded the maximum visible range in each simulated sensor measurement during simulation trials of our planner and compared those with the maximum visible ranges observed by the baseline planner. In every case, our method was able to view a greater distance ahead, resulting in the ability to plan higher-speed actions.

Figures 5a-5c show that our method gives rise to the intuitive behavior of swinging wide around corners to obtain more advantageous views of free space enabling greater stopping distances. Compared to the default baseline behavior pictured in Figure 1, this behavior is qualitatively different and faster.



(a) (b)

Fig. 6: (a) Experimental trajectory of our planner through a hallway in the MIT Stata Center with numbered turns (inset) and expanded view of our planner's trajectory through turn #3 (main image), exhibiting good visibility around the corner as a result of its wide-turn behavior. (b) Our autonomous RC car in turn #2 of the experimental environment.

3) Experimental Demonstration on Autonomous RC Car: We tested our planner in a hallway environment in the MIT Stata Center on the autonomous RC car pictured in Figure 6. The car is equipped with a Hokuyo UTM-30LX LIDAR, which has an artificially limited field-of-view for the purpose of map building.

We picked a section of hallway consisting of three 90° turns, with hallway widths of approximately 1.9m, roughly matching our simulated hallway distribution. In this environment, we conducted 10 trials each for the baseline planner and our planner, using both 60° and 90° fields-of-view. Figure 6 illustrates our experimental map and autonomous RC car.

TABLE I: Experimental results.

FOV	Planner	Success	Mean Speed (m/s)	Max Speed (m/s)
60°	Baseline	0/10	2.25	4.29
	Learned	**8/10**	**3.08**	**5.35**
90°	Baseline	4/10	3.06	5.30
	Learned	**10/10**	**3.43**	**5.55**

Table I summarizes the experimental results. While we have omitted the discussion of why the baseline planner fails much more frequently than our method due to space constraints, this is an interesting and very useful byproduct of our method, which we discuss in [4]. In speed and success rates, our method significantly outperforms the baseline planner. Figure 6a shows a map and trajectory from one of our planner's successful trials using a 60° field-of-view. The close-up view of the third turn illustrates a characteristic example of the good visibility of free space our planner was able to obtain, hence allowing it to maintain a greater speed around the turn.

III. Learning to Predict Collisions from Camera Images

In this section, we relax the safety constraint with respect to unknown space that was imposed in the previous section, and instead use a learned model of collision probability to plan actions that are *empirically* (though perhaps not formally) safe. In previous years, we showed that this method could result in faster navigation [3], and now we employ it again in the context of camera-based navigation. We assume that the dense geometry of the environment is difficult to reliably estimate from camera images beyond a range of several meters (the

range of the Microsoft Kinect, for example, is advertised as 4m), and so learning to predict collisions from camera images enables us to leverage the appearance of images to plan to a greater range and hence anticipate obstacles at greater distance.

In this section, we define the cost of navigation to be equal to the time spent en route from start to goal, plus a fixed penalty on any collisions that may occur along the way. Let m denote the true, hidden, stationary map of the environment, and $\hat{m}_t$ be a partial estimate of the map built from a SLAM process. The map estimate $\hat{m}_t$ is computed from the complete history of geometric sensor measurements $z_0, \ldots, z_T$ up to time T. The geometric measurements z_t may come from a geometric sensor such as an RGBD camera or from processing stereo or monocular camera images (or even from a simulator during training time). We refer to the raw camera image at time t as i_t. The optimal action at time t is chosen according to the following optimization problem:

$$a_t^* = \underset{a_t \in \mathcal{A}}{\text{argmin}} \; J_a(a_t) + J_c \cdot p(c|\hat{m}_t, i_t, a_t) + h(a_t, \hat{m}_t) \quad (6)$$

$$\text{s.t. } g(\hat{m}_t, a_t) = 0. \quad (7)$$

In this control law, $a_t \in \mathcal{A}$ is a dynamically feasible action, $J_a(a_t)$ returns the time duration of action a_t, and $h(a_t, \hat{m}_t)$ is a heuristic function that estimates the remaining time to navigate from the end of action a_t to the goal. The second term is an expected collision penalty in which $p(c|\hat{m}_t, i_t, a_t)$ is the probability that a collision will inevitably occur *beyond the end* of action a_t, given the current map estimate and image.

Finally, $g(\hat{m}_t, a_t)$ is a collision checking function that returns 1 if action a_t intersects an observed obstacle in the map estimate $\hat{m}_t$, and returns 0 if the action lies entirely within known free space or unobserved space. Note that a conventional collision-checking routine would also classify actions intersecting unknown regions of the map as being in collision as well, in order to formally guarantee safety. However, we consider potential collisions beyond the mapped regions of $\hat{m}_t$ to be the responsibility of the collision prediction term, not the constraint. This formulation will allow us to plan actions whose stopping distance exceeds the available known free space in cases where we trust the collision prediction–a topic that we will describe in detail below.

As stated, this receding-horizon control formulation is straightforward, with one exception: The collision probability $p(c|\hat{m}_t, i_t, a_t)$ is an unknown quantity. In a conventional MDP or POMDP scenario, this probability distribution might be derived from the distribution over environments given in the problem definition. However, since we do not know the distribution over the real-world environments in which we hope to plan, we will approximate the distribution $p(c|\hat{m}_t, i_t, a_t)$ by learning a function from data.

In this work, we employ a neural network to predict collision probability. Therefore, we may be tempted to simply model the collision probability term in equation (6) as:

$$p(c|\hat{m}_t, i_t, a_t) = f_c(c|i_t, a_t), \quad (8)$$

where $f_c(c|i_t, a_t)$ is implemented as a neural network that takes an image and action as input. Since we would like to predict collisions based on images, not geometric maps, we have dropped the dependence on $\hat{m}_t$.

However, neural networks may make erratic predictions when queried with an input unlike their training distribution, so we must be able to detect when an input is novel and revert to a default collision probability estimate that is known to be safe, though perhaps conservative. To do so, we introduce a function $f_n(i_t)$, which returns 1 if image i_t is novel (i.e., unlike the distribution of training images), and returns 0 otherwise. We then model the collision probability distribution as follows:

$$p(c|\hat{m}_t, i_t, a_t) = \begin{cases} f_c(c|i_t, a_t) & \text{if } f_n(i_t) = 0 \\ f_p(c|\hat{m}_t, a_t) & \text{if } f_n(i_t) = 1 \end{cases} \quad (9)$$

In this equation, the function $f_c(c|i_t, a_t)$ is a neural network trained to predict collisions, as in equation (8), and $f_p(c|\hat{m}_t, a_t)$ represents a prior estimate of collision probability based on the geometric map estimate $\hat{m}_t$ that can encourage safe and sensible behavior when the robot is faced with a novel-appearing environment for which f_c is untrained. A reasonable prior might be to limit speed such that the robot can stop within the known free space of map $\hat{m}_t$ if need be, and since our geometric sensing range is limited, this speed may be rather slow. In completely novel-appearing environments, the robot will navigate according to the behavior suggested by the prior f_p, whereas in environments of familiar appearance, the robot will trust its learned collision prediction model f_c. Due to space constraints, we do not discuss learning the model $f_c(c|i_t, a_t)$ here, but we have developed similar tools in [3]. In the next section, we will describe the method we use for building $f_n(i_t)$ to detect novel images.

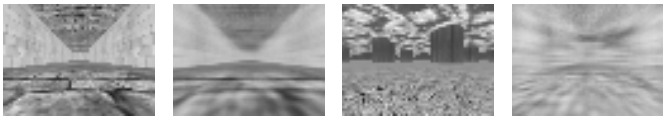
A. Novelty Detection

To detect novelty, we adopt an autoencoder-based approach proposed by Japkowicz et al. [2]. Classification of novel data amounts to a one-class classification problem since we assume that we only have access to the set of images that have been observed and collected by our system. We have no examples of any other "unfamiliar" images. Therefore, we cannot pose novelty classification as a binary classification problem.

To classify images as novel, we use a feedforward autoencoder with three hidden sigmoid layers and a sigmoid output layer. The autoencoder is trained using mini-batch stochastic gradient descent to reconstruct images by minimizing the sum of squared pixel differences between the input and reconstructed output. Let $i^k \in [0, 1]$ denote the intensity value of the kth pixel of image i. The autoencoder loss function is:

$$L_n(i, \hat{i}) = \frac{1}{K} \sum_{k=1}^{K} (i^k - \hat{i}^k)^2, \quad (10)$$

where $\hat{i}$ refers to the output of the autoencoder, i.e., the reconstructed image. The number of pixels in i and $\hat{i}$ is K.



(a) Input. (b) Output. (c) Input. (d) Output.

Fig. 7: (a) A familiar input image of a simulated hallway, drawn from the training distribution, and (b) the associated accurate autoencoder reconstruction. (c) A novel input image of a simulated "cylinder forest," *not* drawn from the training distribution, and (d) the resulting inaccurate reconstruction.

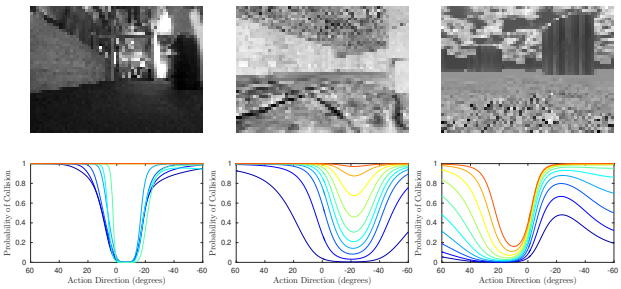
This novelty detection approach assumes that the autoencoder network will learn a compressed representation capturing the essential characteristics of its training dataset such that it will reconstruct similar inputs with low reconstruction error. On the other hand, when the autoencoder is presented with a novel input—that is, an input unlike any of its training data—we expect the reconstruction error to be large. Therefore, to classify novel images, we need only to set a threshold and determine whether the reconstruction error for a given image falls above or below that threshold.

To evaluate this novelty detection approach on a visual navigation domain, we trained an autoencoder on a set of simulated camera images from a hallway-type environment and tested it on simulated images of hallways (the training distribution) as well as forests of cylindrical obstacles (the novel test distribution). We use three hidden layers with sigmoid activations and 200 nodes per layer. For computational efficiency, we use very low resolution 60×44 images in this work, so the input and output layer have dimension 2640.

Figure 7 illustrates two examples of novelty detection. Figure 7a illustrates an input image of a simulated hallway drawn from the distribution of images used to train the autoencoder and 7b shows the corresponding reconstruction with a reconstruction error of $L_n = 4.5\times10^{-3}$, which is lower than the threshold of 6.3×10^{-3}, resulting in a classification of "not novel". On the other hand, Figure 7c illustrates an input image of a simulated forest of cylindrical obstacles *not* drawn from the training distribution and 7d shows the corresponding reconstruction, which has a reconstruction error of $L_n = 4.4\times10^{-2}$, which is much greater than the threshold, resulting in a classification of this image as "novel". It is visually apparent that the reconstruction in Figure 7b is much more accurate than that in Figure 7d, intuitively indicating that the simulated "cylinder-forest" image does not follow the same encoding that the autoencoder learned from the simulated hallway environment.

B. Vision-based Collision-Prediction Results

We tested our approach in experiments using a small autonomous RC car. We initialized the system with an empty dataset and a randomly initialized collision prediction network and autoencoder. The autoencoder threshold was initialized to zero so that all images appeared novel. We then repeatedly deployed the robot to drive to a specified goal location, collecting images during each run, labeling them, and periodically re-training the collision predictor and autoencoder networks on the growing dataset.



(a) Real hallway. (b) Simulated hallway. (c) Simulated forest.

Fig. 8: Illustration of predicted collision probabilities. Each image is illustrated with the learned collision probability model below it, with the horizontal axis representing action direction and color representing speed at the end of the action, from 1 m/s to 10 m/s.

In our experimental results, we performed trials using two different sensor configurations. In one configuration, we used monocular visual-inertial odometry based on GTSAM for state estimation, and an Xtion RGBD camera for depth estimation. In the other sensor configuration, we used a Hokuyo UTM-30LX planar LIDAR scanner to perform scan-matching for state estimation, and (artificially limited to 110° field-of-view and 5 m range) for geometric perception. In both experimental configurations, we used a Point Grey Chameleon camera with a 110° horizontal field-of-view as input to our collision-prediction and novelty-detection models.

While this work was targeted at camera-only perception, the drift and error in visual state estimation during aggressive driving prevented us from achieving maximum performance and reliability in the LIDAR-free configuration. Therefore, we employed the LIDAR purely to test the merits of our learning approach under the idealized circumstances of highly accurate state estimation. We emphasize that in all tests, regardless of the sensor suite being used for state estimation and mapping, collision prediction is performed only on *camera images* combined with the action description, as described above. Therefore our results will easily extend to future implementations of camera-only navigation as visual state estimation and depth estimation become more reliable.

1) Visualizing Neural Network Predictions: To illustrate the effectiveness of collision prediction using camera images, Figure 8 shows the predictions of collision probability on two simulated images and one real camera image encountered during an experimental run. In each case, these networks had been trained with a large dataset of relevant training images, so novelty is not a factor.

The graph below each image illustrates the collision probability as a function of the steering direction of the action (i.e., $\text{atan2}(a_y, a_x)$), shown on the horizontal axis, and the speed at the end of the action, which is indicated by the color of each line. The speeds illustrated vary from 1 m/s (blue) to 10 m/s (red), and intuitively illustrate that faster actions carry a



(a) LIDAR-mapped hallway environment for training and testing.



(b) Xtion-mapped novel environment used for LIDAR-free testing.

Fig. 9: Maps of experimental environments.

greater probability of resulting in future collision.

Figure 8a illustrates a case where the robot is approximately aligned with the axis of a real hallway, and therefore the learned model assigns low probability of collision to actions aligned with that axis. Figure 8b illustrates the robot entering a turn in a simulated hallway, so the probability model assigns lowest probability of collision to actions that bear 20 degrees to the right. Finally, Figure 8c illustrates the robot in a simulated forest of cylindrical obstacles. Since there is a long gap between obstacles just to the left of center of the robot, the probability of collision is lowest at this point. Conversely, since there is an obstacle to the right of the image, the predicted collision probability peaks at an angle of 20 degrees to the right. However, for low speeds, the probability drops off again for actions that bear even farther right, since the robot could drive around the right side of the obstacle.

2) Experimental Results using LIDAR: We performed a total of 80 experimental trials using a small autonomous RC car, retraining every 10 runs (approximately every 1000 images). A generated map of the experimental environment is shown in Figure 9a and an image from this environment is shown in Figure 11a.

As the system collects more training images, the fraction of images considered novel decreases, resulting in a corresponding increase in speed of navigation, as illustrated in Figure 10. We observe an average speed increase of 25%, from 3.38 m/s to 4.24 m/s. Since the experimental environment is small compared to the simulation environment, a substantial portion of the navigation time is spent accelerating from the start and decelerating at the goal. Therefore, we also provide the average maximum speeds observed for each amount of training data, where we see a much more substantial improvement of 55%, from 3.79 m/s to 5.90 m/s. This improvement in maximum speed implies that if the experimental environment were larger, we would observe a greater improvement in average speed as well.

Figure 12 illustrates the autoencoder reconstruction error for three categories of data after training on nearly 9000 training images from numerous runs through the experimental environment depicted in Figure 9a and 11a. The blue histogram and PDF indicates the reconstruction error of the training

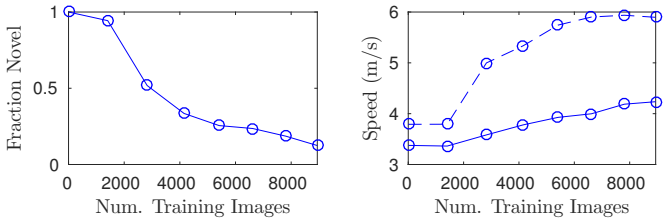
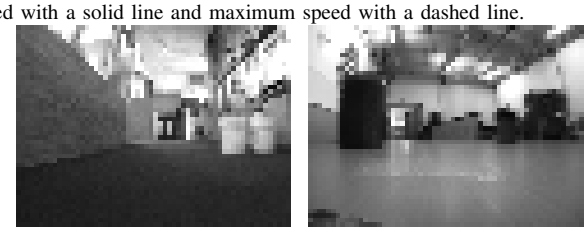


Fig. 10: Experimental results. In the right panel, we illustrate the average speed with a solid line and maximum speed with a dashed line.



(a) Familiar Environment. (b) Novel Environment.

Fig. 11: Maps of experimental environments.

images, from which the threshold (dashed line) is computed. The red histogram and PDF represents the reconstruction error of approximately 1000 images collected during testing runs in the same environment, but not used for training the autoencoder. While the reconstruction error of these images is not as low as the training images, most fall below the novelty threshold. Finally, the black histogram and PDF illustrate the reconstruction error of a truly novel environment–a well-lit lab area with obstacles (Figures 9b and 11b). Visually, this environment is quite different from the training environment, and that is reflected in the much greater reconstruction error, with all images falling above the novelty threshold.

3) Experimental Results using Monocular VIO and RGBD: We performed experiments on our autonomous RC car using a LIDAR-free sensor configuration in order to demonstrate steps toward a completely vision-based learning and navigation pipeline. However, as noted previously, inaccuracies in the monocular VIO state estimate made this sensor configuration unsuitable for demonstrating the entire training and testing procedure. Nevertheless, we were able to test the LIDAR-free sensor configuration in both environments depicted in Figures 9 and 11 using the collision-prediction and autoencoder models previously trained with the other sensor configuration.

In the hallway training environment, we achieved a mean

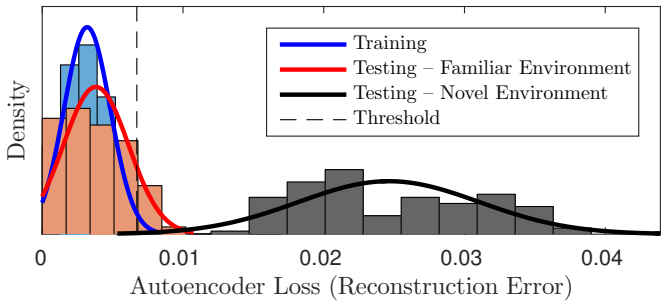


Fig. 12: Histogram of autoencoder reconstruction errors for experimental data.

speed of 3.26 m/s and a top speed over 5.03 m/s. This result significantly exceeds the maximum speeds achieved when driving in this environment under the prior estimate of collision probability before performing any learning. On the other hand, in the novel environment, for which our model was untrained, the novelty detector correctly identified every image as being unfamiliar. In the novel environment, we achieved a mean speed of 2.49 m/s and a maximum speed of 3.17 m/s.

We emphasize that these results represent an improvement of 31% in mean speed and 59% in maximum speed after learning, for monocular vision-based navigation using a neural network for planning. Naturally, the robot navigates more slowly in the novel environment than it does in the familiar one, but our method enables the robot to safely navigate in the novel environment, potentially gathering data and retraining, whereas that would not be safely feasible without our novelty-detection approach.

IV. Robust 3D Perception

The previous section showed an approach to learning to use monocular images to predict collision probabilities beyond the lookahead capability of two types of active depth sensors, but still used the active depth sensors for state estimation and for collision checking using traditional 3D world models, within the range of the active depth sensors. In this section, we describe how we incorporated depth perception and image-based collision checking with a new method for image-space temporal fusion of depth maps, to improve the robustness of image-space motion planning without the use of traditional 3D world models.

Depth perception over wide horizontal and vertical fields of view is more powerful for collision prediction and evaluation of candidate trajectories than using just monocular images or single-axis lidar scanners like the Hokuyo. How to achieve such depth perception reliably, with adequate resolution and maximum range, and with compact, lightweight, low power sensing and computing hardware is a challenging and complex problem. We have discussed this at length in [8], where we introduced the use of inverse range images from stereo vision and image-based configuration space (c-space) expansion for collision checking of candidate trajectories for micro air vehicles. In [9] we extended this from using just a forward-looking stereo camera pair, with a corresponding forward-looking, planar image space depth representation, to using a 360-degree cylindrical depth map centered on the robot as a persistent representation of the nearest surface in all directions. We also added side-looking cameras with a visual structure-from-motion (SfM) algorithm to provide depth perception over a wider field of regard.

These approaches are computationally cheaper than maintaining and doing collision checking with 3D world models like occupancy grids or voxel maps. We have demonstrated successful use of these approaches in outdoor operation in areas with a high degree of visual texture (e.g. forests). These approaches can also provide depth data to augment image data for learning to assess the utility of trajectories. However, there

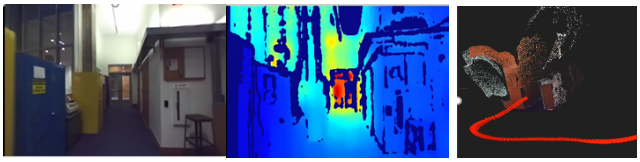


Fig. 13: Stereo depth and motion estimation results with a passive Zed Stereo camera. (a) Sample image of the navigation environment, (b) generated depth map, and (c) state-estimation of the path traveled using stereo visual odometry

are still three key limitations with this. First, as a practical matter, until very recently the computational power available on small robots only permitted depth perception with stereo vision and SfM with relatively low image resolutions and frame rates, which limits lookahead distance, reaction time, and therefore limits robot velocity. Second, in areas that lack visual texture, passive range sensors (stereo and SfM) are forced to use model-based assumptions to hypothesize the range to textureless surfaces, which is error-prone. Third, all range sensors are subject to range errors with some probability. Occupancy grids and voxel maps address this by temporal fusion of successive frames of range data, but our initial work with image space representations did not have such a mechanism.

We have now addressed these limitations by incorporating approaches to real-time, hybrid active/passive depth sensing (Section IV-A) and developing a novel method for temporal fusion of depth maps in image space (Section IV-B).

A. Evaluating Approaches to Hybrid Active/Passive Depth Perception

Combinations of active and passive depth perception have potential benefits of improved depth perception on featureless surfaces and in low light (with active depth sensors), while providing potential lookahead distance advantages (with passive sensors). In FY16 we examined two relatively new classes of active depth sensor that are now providing high frame rates in very compact packages. One class is based time of flight, AM-CW flash lidars based on detectors developed by PMDTec, which are now appearing in some smartphones. These are currently limited to maximum ranges of a few meters indoors by eye safety limits on the optical emitters, and to much shorter ranges in sunlight. The second class is based on ”active stereo”, which uses near infrared illuminators to project texture patterns on the scene and measures range by stereo triangulation with two cameras; the Intel RealSense sensors are in this class. Initial versions of active stereo systems use cameras only sensitive in the near infrared, so they also had very limited range indoors under fluorescent lighting. Future versions will use broad band cameras, which should give them better range in ambient lighting indoors and outdoors, combining the advantages of active and passive ranging in one device.

The commercial sensors above provide high frame rates with very compact hardware. PMDTec-based flash lidars are still fairly low resolution (e.g. 320x240 pixels), but active

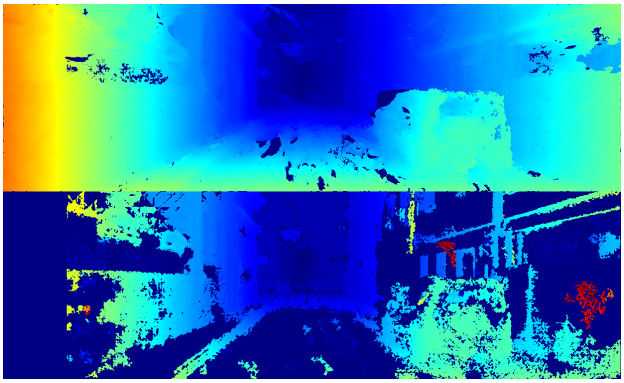


Fig. 14: A depth map created by GMM-based image space temporal fusion (top) and one of the original depth maps (bottom). The fused map has filled in gaps and eliminated several patches of erroneous depth data.

stereo systems now provide at least VGA resolution (640x480 pixels). Unfortunately, their range limitations strongly constrain their utility in our application. Therefore, to provide fast, high resolution, wide field of view depth data from images in this last year of the program, we also examined options for higher performance onboard processors for stereo vision. The most readily available approach uses the NVIDIA Jetson TX1 processor board with a commercially available stereo camera head and TX1 GPU-based stereo vision software library from Zed Stereo. These cameras provide a 110 degree field of view and VGA resolution depth maps at up to 100 frames/sec; the software library also provides stereo vision-based visual odometry, which gives 6-DOF robot motion directly. We evaluated the depth maps and visual odometry from this system with data from MIT and found it to be quite promising (Figure 13).

A potential drawback of this solution is that the cameras have a rolling shutter readout architecture, which can induce motion artifacts in depth maps under high speed robot motion. We have evaluated this qualitatively with stereo image data sets acquired by manually carrying the cameras at speeds up to about 3 m/s, with satisfactory results.

B. Temporal Fusion of Depth Maps using GMMs

Occupancy grids and voxel maps inherently provide temporal fusion of successive frames of range data, which progressively fills in gaps in individual frames and reduces the probability of erroneous range measurements. Our initial approach to collision checking with inverse range images did not provide such a mechanism. This made it vulnerable to false alarm obstacles or missed obstacle detections. To address this, we developed a novel approach to depth map fusion in image space that uses Gaussian Mixture Models (GMMs) to track multiple depth hypotheses at each pixel. The environment is modeled from an ego-centric point of view, where each pixel is represented by a mixture of Gaussian inverse-depth models. The image-based GMM representation is propagated between time steps using estimates of the camera motion, then updated with the next frame of depth data. Details of the algorithm are described in [11] and sample results are shown in Figure 14. We have evaluated this approach quantitatively with public data sets, integrated tested it into our real-time depth perception pipeline, and tested it in outdoor operation in challenging scenes [10].

V. Semi-supervised Near-to-Far Learning

Previous sections used geometry-based 3D perception methods that provide robust and reliable scene estimation at near-range. In order to overcome the inherently myopic nature of these geometric approaches, we developed a monocular, end-to-end learning-based approach that enables planning beyond the range of 3D depth sensors. However, such learning based methods are limited by the large amounts of labeled data needed to train these systems, as well as their reliability in new and dynamically changing environments. This motivates the need to develop systems that alleviate these problems by exploiting the complimentary strengths of both geometry and data-driven learning methods. Here, we take a step in this direction with a framework for adaptive perception using near-field depth perception for self-supervised online learning.

Continuous online learning of scene segmentation would allow a system to constantly adapt to its local environment and avoid the need to learn a universal scene classifier. This requires a method to automatically generate large amounts of training labels in real-time. We achieve this by using near-field range data to generate segmentation labels in real-time, which are then used to adaptively train a long-range classifier online. Learning from such self-supervision automates labeling, extrapolates scene knowledge from short range to long range, and is referred to as near-to-far learning.

A. Scene Segmentation

Our system takes as input an RGB image, a depth map, which can be obtained either using an active or passive range sensor, and the estimated pose of the vehicle. The problem of learning scene segmentation is then defined as associating each pixel of an input image to a semantic class. In this work, we reduce the semantic scene understanding problem to learning a two-class classification: obstacle and free-space. This is the minimum classification required by an autonomous navigation system onboard an MAV to plan collision-free trajectories.

Recently, neural networks with deep learning have achieved state-of-the-art performance in the task of semantic scene segmentation. Inspired by these results, we approach the problem with a CNN that can be trained end-to-end to predict a map of class labels. Our architecture is based on the recently introduced Fully Convolutional Networks (FCNs) [12].The general structure of such networks consists of an encoder (or contraction) segment and a decoder (or expansion) segment. The contraction segment is typically a classification network, where the fully-connected layers have been substituted with convolutional layers; this adaptation transforms the network into a deep filter that preserves spatial information, since it only consists of filtering layers that are invariant to translation. The encoder maps the input to a low resolution representation,

while the expansion maps the low resolution feature maps to high-resolution segmentation output.

We represent our network model as $f(x, \gamma)$, that maps an input image x to the target segmentation. The model is described by the network parameters γ and is learned online by minimizing its error output for an instance x_i given an output ground-truth label y_i:

$$\gamma = argmin_\gamma \sum_{i=1}^{N} L(f(x_i, \gamma), y_i) \tag{11}$$

where N is the self-supervised training set generated from near-range geometry-based segmentation, and L is the cross-entropy (*softmax*) loss. During inference, the *softmax* is replaced by an *argmax* function, so as to provide a single output per pixel.

B. Self-supervised Training from Geometry

We segment the near-range 3D information into obstacle and free-space regions based on the typical application of ground-plane estimation. Ground-plane estimation typically entails finding the largest planar region in the 3D scene. We assume the free-space in front of the MAV is locally flat and that this plane corresponds locally to the ground. Any points that are outliers with respect to the current ground plane model are then distinguished and assumed to be obstacles.

Given a point x_w on the ground plane in the Euclidean coordinates, referenced to an arbitrary frame of reference, a plane in this space can be parametrized by a vector $c_w^T x_w = 1$. It can be shown that under the projective transformation of a pinhole camera model, every such ground plane corresponds to a linear model in image coordinates and disparity, and can be written as:

$$\alpha_0 x + \alpha_1 y + \alpha_3 = d \tag{12}$$

where (x, y) are the image coordinates of the pixel with the observed disparity d. Given enough points on the ground, it is possible to estimate the model parameters in a robust least squares sense using RANSAC.

Once the ground plane has been estimated using the above described method, we use it to segment each pixel for which there is valid depth information as either obstacle or free-space. This is accomplished using a simple threshold on the residual of the fit for each observation, i.e. the ith pixel is labeled ground if $|\alpha_1 x_i + \alpha_2 y_i + \alpha_3 - d_i| \leq d_g$ (where d_g is a constant threshold), and is labeled an obstacle if the residual exceeds another threshold. This near-range classification is then used as a self-supervised training set for the learning algorithm.

Training a neural network typically requires a large amount of data samples for convergence of the parameters. However, several strategies have been adopted in the literature to handle this. One effective method initializes the network parameters with a network trained on a larger, generic dataset in an unsupervised manner; to apply it to a new task, one can fine-tune the last layer of the net, which provides class confidences, based on the required task. This exploits the observation that these pre-trained networks are a compact and yet rich representation of the images in general, since they are trained extensively on large visual datasets.

We perform training in a similar multi-stage fashion. In the offline step, we initialized the decoder part of the network with the AlexNet architecture. The encoder part is initialized using xavier initialization, and fine-tuned online in real-time by back-propagation using a stochastic gradient descent in a sliding window fashion, where each minibatch consists of the last N frames.

To improve the runtime efficiency of the network, we modified the network hyper-parameters; reduced the padding of the first convolutional layer from 100 to 10 pixels (to enable faster training), used a higher learning rate, from $1e^{-10}$ to $1e^{-9}$ and lower momentum of 0.90 instead of 0.99. We also changed the fixed learning rate (Lr) with a poly-learning policy

$$L_r = L(1 - i/max_i)^p \tag{13}$$

where L is the base learning rate, i is the learning step and p is the power index. These modifications make our approach more run-time efficient without significantly affecting classification error rates; quantitative results are reported elsewhere [13].

C. Performance Evaluation

We analyzed the performance of this method and demonstrate its efficacy through quantitative evaluate on benchmark data sets [13] and with real-world experiments using a modified version of the Asctec Hummingbird quad-copter platform. Our flight computer is an NVIDIA Jetson TX1 board, which incorporates a quad-core ARM processor and embedded GPU with 256 Cuda cores. Onboard, we have a downward looking camera for state estimation a and front-facing Intel RealSense R200 structured light sensor, which provides visual and range data. We use the Tensorflow deep learning library with CuDNN backend for our implementation, and all the processing is done onboard.

We demonstrated the efficacy of our proposed approach through real-world flight tests in a varied set of indoor and outdoor environments, ranging from untextured building corridors to natural cluttered scenes. During the experiments, the MAV was manually flown by a pilot, while the near-to-far segmentation module was run in real-time as software-in-the-loop. We decoupled the segmentation module from the control and planning layer, so as to independently evaluate the performance of the proposed approach. Figure 16 shows the qualitative results from some of the flight experiments.

It can be seen that self-supervised learning in a continuous fashion performs better, especially in the scenarios where the environment changes dynamically. In the first example, we notice that as soon as the person appears in near-field, the perception algorithm is able to adapt its learning to classify people as potential obstacles; this is completely missed in the supervised scenario where the training set does not have humans as an annotated class. Similar behavior is observed in outdoor environments, in the context of narrow tree trunks

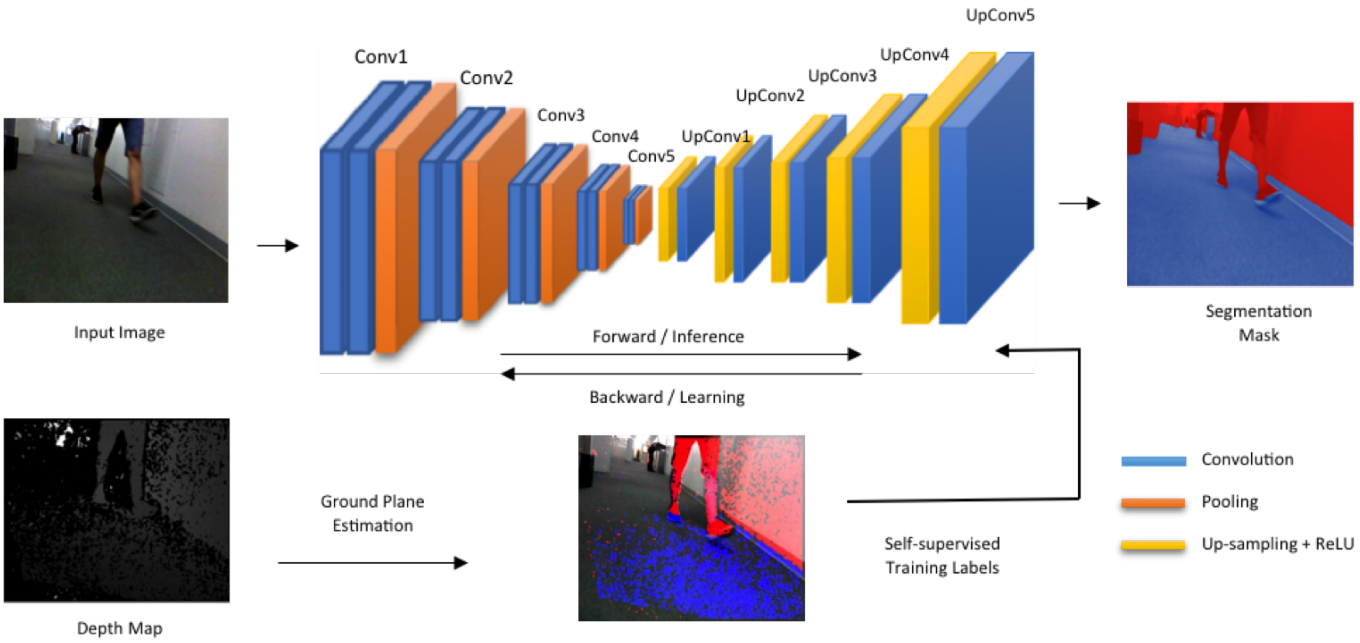


Fig. 15: Schematic overview of the near-to-far scene segmentation approach. Geometric information from the near-range is used to generate self-supervised training labels in real-time via ground plane estimation. These labels are then used to adaptively train a classifier for appearance-based scene segmentation.

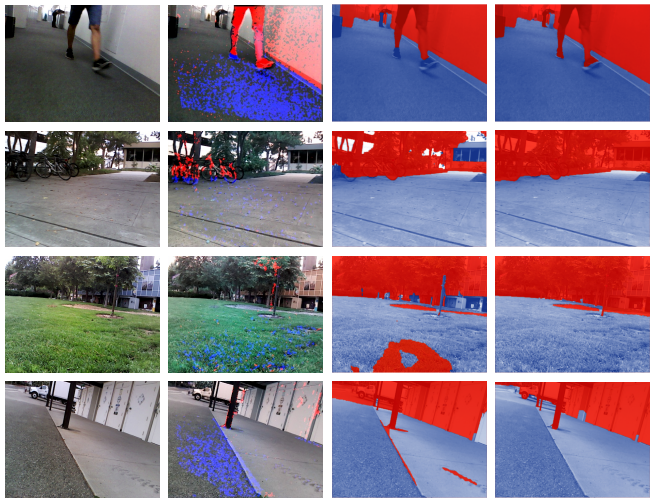


Fig. 16: Qualitative results of near-to-far scene segmentation from MAV flight tests in different indoor and outdoor environments. (Col 1) Input RGB image. (Col 2) Near-range geometry-based segmentation overlayed on the input image; this is further used for online training of the classifier. (Col 3) Overlayed segmentation mask for the baseline algorithm (without online learning component). (Col 4) Overlayed segmentation mask for the proposed approach. Note: Red is obstacle, blue is free-space.

as shown in the third row. Further, the ability of the near-to-far scene segmentation to learn from its current 'local' information is evident in the last example where it can successfully segment the planar surfaces, even though they have a considerably different texture and visual appearance.

Computationally, our geometry-based near-range segmentation module runs at $10hz$, and the near-to-far scene segmentation module runs at $1Hz$; this includes both the inference and the online update step. These run-times are sufficient for our application of high-speed autonomous navigation, where we envision a slower long-range strategic planning layer running in tandem with a faster near-range geometry-based reactive avoidance system. Closing the loop with planning and control to show system-wide advantage of the proposed algorithm is beyond the scope of this effort, and will be addressed in future work.

VI. Conclusions and Future Work

In FY16 we developed a learning based method for predicting collision probabilities directly from monocular camera images, effectively increasing the planning horizon beyond the limited range of environment geometry that can be reliably inferred from range sensors. In the final year of this project, we completed a high speed depth perception pipeline that combines frame rate depth perception, a novel approach to image-based temporal fusion of depth maps, and visual motion estimation as an alternative to the use of Hokuyo lidar and occupancy grid. We also developed a framework for using depth estimation to segment nearby portions of the scene in ground and obstacles, which automatically provides training labels for a near-to-far learning approach to online update of a CNN-based terrain classifier. This essentially extrapolates the near-field geometric knowledge from the depth sensor to far-field semantic scene understanding. Although integrating this perception pipeline with motion planning was not completed in this work, we expect that this approach will enable a very efficient generalization of our planning work that increases confidence in collision predictions at long range by using

inferences about ground versus obstacle semantic labels that are learned from near field depth data.

REFERENCES

[1] T. Fraichard and H. Asama. Inevitable collision states–a step toward safer robots? *Advanced Robotics*, 18(10):1001–1024, 2004.

[2] N. Japkowicz, C. Myers, M. Gluck. A novelty detection approach to classification. In *International Joint Conference on Artificial Intelligence*, 1995.

[3] C. Richter, W. Vega-Brown, and N. Roy. Bayesian learning for high-speed navigation in unknown environments. In *Proc. ISRR*, 2015.

[4] C. Richter and N. Roy. Learning to plan for visibility in navigation of unknown environments In *Proc. ISER*, 2016.

[5] C. Richter, W. Vega-Brown, and N. Roy. Markov chain hallway and Poisson forest environment generating distributions. Technical Report MIT-CSAIL-TR-2015-014, 2015.

[6] A. Bachrach et al. Estimation, planning, and mapping for autonomous flight using an RGB-D camera in GPS-denied environments. *International Journal of Robotics Research*, 31(11): 1320-1343, 2012.

[7] M. Otte et al. Path planning in image space for autonomous robot navigation in unstructured outdoor environments *Journal of Field Robotics*, 26(2):212-240, 2009.

[8] L. Matthies et al. Stereo vision-based obstacle avoidance for micro air vehicles using disparity space In *IEEE International Conference on Robotics and Automation (ICRA)*, 2014

[9] R. Brockers, A. Fragoso, B. Rothrock, C. Lee, and L. Matthies. Vision-based obstacle avoidance for micro air vehicles using an egocylindrical depth map. In *International Symposium on Experimental Robotics (ISER)*, 2016.

[10] A. Fragoso, C. Cigla, R. Brockers, and L. Matthies. Dynamically feasible motion planning for micro air vehicles using an egocylinder. In *11th Conference on Field and Service Robotics (FSR)*, 2017.

[11] C. Cigla, R. Brockers, and L. Matthies. Gaussian mixture models for temporal depth fusion In *Winter Conference on Applications of Computer Vision (WACV)*, 2017

[12] J. Long, E. Shelhamer and T. Darrell. Fully convolutional networks for semantic segmentation In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015

[13] S. Daftry, Y. Agarwal, and L. Matthies. Learning near-to-far scene segmentation for micro air vehicles. In *(under review)*, 2018

5. Joint Experiments (JX) Thrust

Thrust Leader: L. Matthies, JPL

ARL Technical Leader: David Baran, Computational and Information Sciences Directorate

5.1 Introduction

The JX thrust serves as a focal point to integrate and evaluate technical capabilities on platforms performing elements of mission-like scenarios, through a set of research tasks that involve collaboration between multiple centers and with ARL. Research tasks in the JX thrust have evolved through four stages during the MAST program. The current focus is on increasing the speed of autonomous navigation among static obstacles; beginning to enable MAV navigation among moving obstacles; autonomous day/night MAV landing on elevated structures for perch-and-stare scenarios; distributed, SLAM for large, heterogeneous robot teams; and transferring autonomous behaviors that were developed on large (>1 kg) MAVs to smaller platforms (<250 kg). This report provides a historical perspective on the evolution of the JX thrust through all four stages and summarizes recent progress in the current focus areas.

5.1.1 Research Vision

The 3 core research thrusts (MCE, SPP, and CNC) address capability gaps at the level of component technologies. The JX thrust serves as a focal point to integrate and evaluate technical capabilities on platforms performing elements of mission-like scenarios, through a set of research tasks that involve collaboration between multiple centers and between the consortium and ARL. This also facilitates technology transition to other government programs and the private sector.

The technologies and capabilities addressed by the JX thrust have evolved throughout the program, as capabilities available from the core research tasks have evolved and matured and as tasks within the JX thrust achieved their objectives. Key themes that emerged early in the program and that remain important elements of the JX thrust include autonomous navigation among static and moving obstacles; autonomous landing of MAVs on elevated structures to enable perch-and-stare missions; and distributed SLAM by heterogeneous, air-ground teams of robots. The focus in the final years of the program also includes transferring autonomous navigation capabilities originally developed for MAVs weighing over 1 kg to smaller MAVs weighing 150 to 250 g.

5.1.2 Army Relevance and Technical Potential

The research vision of the JX thrust is aligned well the US Army Robotic and Autonomous Systems Strategy (RAS). The technologies developed and demonstrated here address a key need of small units and dismounted warfighters for improving reconnaissance while providing greater standoff from threats and reducing the burden of equipment to be carried. Successes in this program have led to transitions to other government programs and helped seed the rapidly growing private sector interest and progress in small, autonomous drones.

5.2 Technical Challenges: 2008–Present

JX tasks were first defined for the FY10 Annual Program Plan (APP) and were formalized as a thrust area in FY11 under the leadership of the Integration center. The JX thrust has evolved through four stages, as summarized in the following subsections. Highlights of current work on tasks in the fourth stage are summarized in Section 5.3.

5.2.1 Stage 1: FY10, Pairwise Component Integration Tasks

JX tasks were defined initially to stimulate collaboration within and among the four MAST technical centers and with ARL, and to bring component technologies from different research efforts together to explore issues that arose as the level of integration increased; progressively more integrated and sophisticated experiments were expected in successive years of the program. Specific tasks were defined in a bottom-up fashion by teams of investigators who felt they could productively collaborate in the spirit of JX goals. A total of 10 tasks were defined, grouped into 3 broad areas:

1) Advances in small, agile, air and ground vehicles for exploring unstructured 3-D environments. A set of 3 tasks addressed research issues involving behavior of a single robot at a time:

    i. Gust studies: characterized and understood the effect of wind and gusts on the flight performance of MAST-scale aircraft, and developed new approaches to enable improved flight in such conditions (Mechanics and Integration centers).

    ii. Sensor fusion for autonomous navigation: evaluated obstacle avoidance performance for small vehicles using 1) sonar and optic flow and 2) sonar and LIDAR (Mechanics and Processing centers).

    iii. Aerial deployment of sensor network nodes: developed and evaluated a system using a quadrotor to deploy multiple mesh

network nodes to receive surveillance video through the network (Mechanics and Integration centers).

2) Exploiting novel, lightweight electronics and sensors for autonomous navigation and behaviors. Five tasks focused on integration and evaluation of progress on sensor, electronics, and/or perception algorithms for autonomous navigation and reconnaissance:

   i. Autonomous landing: used onboard monocular SfM-based depth perception to detect and land on elevated, planar structures (e.g., small-scale, laboratory surrogates for buildings) for persistent perch-and-stare type missions scenarios. The aircraft motion was estimated with external motion capture cameras (Vicon system) (Processing and Integration centers).

   ii. Autonomous ingress: used onboard monocular SfM to detect and fly through open windows with MAVs, with Vicon providing aircraft motion estimates (Processing and Integration centers).

   iii. Platform-integrated antennas: integrated antennas into wings of a flapping air vehicle and into the body of a crawling ground vehicle, and evaluated communication performance (Microelectronics and Mechanics centers).

   iv. Integrated point-sampling sensors: integrated small reconnaissance-oriented sensors (gas, radiation, and airflow) developed under MAST into small robots, then evaluated performance of the sensors and robot reconnaissance behaviors (Microelectronics, Processing, and Mechanics centers).

   v. Navigation with radar: developed a simulation of MAV navigation with the 235-GHz radar under development in MAST and used the simulation to compare performance with only optic flow sensing, only radar sensing, and with both (Microelectronics and Mechanics centers).

3) Assessing performance and operational capabilities of collaborative ensembles of MAST platforms. Two tasks focused on teams of robots:

   i. Coordinated SLAM: evaluated multi-robot SLAM in an indoor, building search scenario, with limitations on communications and first steps toward using semantic object recognition (of doorways) to reduce the computational cost of mapping (Processing and Integration centers and ARL).

ii. Robot ensemble communication with repeaters: evaluated small communication repeaters under development in MAST for multi-robot reconnaissance scenarios (all 4 centers and ARL).

Figure 5-1 illustrates results from two of these tasks. Research on many of the others is discussed further and illustrated in subsequent sections.

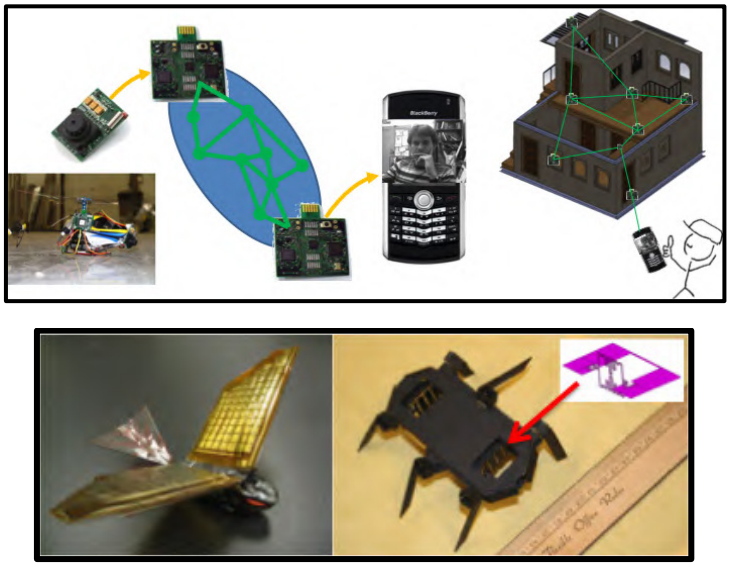


Fig. 5-1 Sample results from FY10. (Top) Architecture of a demonstration of multi-hop video transmission over small, low-bandwidth, mesh network nodes that can be deployed by MAVs.[1] (Bottom) Antennas integrated into the structure of a wing with a solar array[2,3] and into a small, crawling ground vehicle.[4]

5.2.2 Stage 2: FY11–12, JX Thrust Established with 4 Focus Areas Leading to Year 5 Capstone

The JX thrust was formalized under the Integration center in FY11. Research themes were structured into four tasks with distinct themes and with research goals driven by the three mission scenarios motivating the program: small-unit search of intact buildings, small-unit search in caves or demolished buildings, and small-unit perimeter defense. Results over FY11–12 facilitated technology transition and led to capstone demonstrations in the last year of the base period of the program (FY12). The tasks, participating centers, and main research objectives were the following:

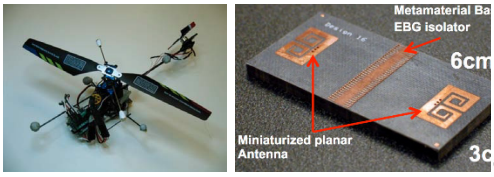
1) ***Ground vehicle capabilities (Mechanics, Processing, and Integration centers):*** This task expanded the theme of FY10 Task 1.2 by using compact sensors and a 2-g, 3-W OMAP3 system-on-a-chip processor module to develop and evaluate autonomous navigation behaviors for MAST-scale ground vehicles. Behaviors included the following:

- Enabling the OctoRoACH vehicle developed by the Mechanics center to follow humans who were tracked using an onboard array of pyrometer sensors;
- As a potential transition activity, using the OMAP3 processor in the FirstLook commercial throwbot from iRobot to do autonomous visual servoing toward user-designated objects in images.

2) ***Air vehicle capabilities (Mechanics, Processing, and Integration centers):*** This task extended FY10 work on behaviors for individual air vehicles in several directions:
   - Further quantitative characterization of MAV flight control in gusts, using a programmable turbulence generator, a Vicon system for accurate motion estimation, and experimenting with H-infinity control to improve position and attitude control.
   - Integration of onboard visual-inertial SLAM algorithms using an onboard camera and IMU to provide self-contained state estimation, and using this to reproduce the FY10 autonomous landing experiments without reliance on a Vicon system.
3) ***Team communication (Processing and Microelectronics centers and ARL):*** This extended FY 10 Task 3.2 by the following:
   - Integrating improved communication repeater modules from the Microelectronics center into ground robots in the Processing center.
   - Developing algorithms to use received signal strength indicators (RSSIs) to estimate the relative location of robots indoors.
   - Conducting experiments on communication and localization performance during indoor, multi-robot operation.
4) ***Air-ground collaboration (Processing, Mechanics, and Microelectronics centers):*** This expanded the theme of FY10 Task 3.1 on multi-robot collaborative behaviors in several directions:
   - Using an air vehicle to deploy a small ground vehicle for persistent surveillance, and using video from the air vehicle to enable visual servo-based guidance of the ground vehicle to desired locations.
   - Innovations in cooperative, multi-robot mapping with centralized mapping algorithms.

- Integrating a micro gas chromatograph sensor from the Electronics center into a ground vehicle developed in the Processing center and conducting experiments with autonomous mapping of chemical gradients.
- Developing algorithms to multiple air vehicles to conduct periodic reconnaissance sorties for repeated observation of the same location.

MAV flight in gusts was an ongoing research theme in the JX thrust. Initial work from FY10 on characterizing the effect of gusts on MAVs was extended in FY11–12 to conduct more quantitative evaluation in a variety of conditions[5] and develop improved control algorithms for flight in gusts.[6] Communication for small robot teams was also an ongoing theme; in FY11–12 the focus in the JX thrust shifted from the mesh network nodes of FY10 to integrating and evaluating the utility of small radio repeaters in teams of ground robots. These two examples are illustrated in Fig. 5-2. Research on team mapping, autonomous landing for perch-and-stare, and other topics expanded in the option years of the program.



Controller	Vehicle	CEP (cm)	RMS (m)	Max Excursion (m)
PD	Flybar	9	12	23
H_∞	Flybar	13	12	21
PD	Flybarless	7	8	15
H_∞	Flybarless	3	4	9

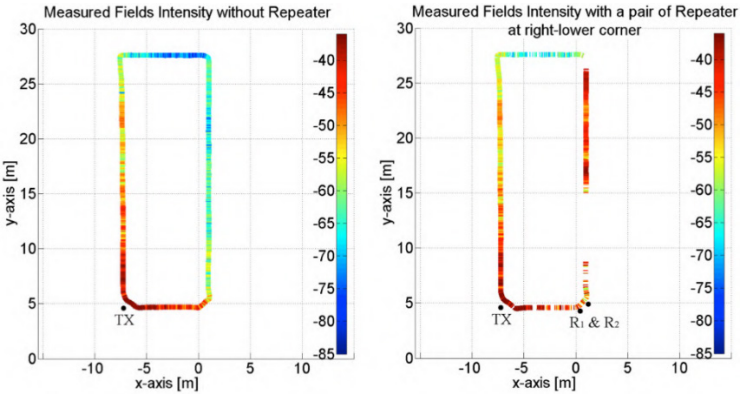


Fig. 5-2 Sample results from FY11–12. (Left) Photo of a small helicopter with a flybar; flight performance was compared (in the table) to one without a flybar, with either PD or H-infinity control.[6] Also shown: a small wireless radio repeater used for robot team communication experiments. (Right) Results of experiments communication signal strength in corridor loop with and without radio repeaters on small ground robots.

5.2.3 Stage 3: FY13–14, Option Years with New Tasks Selected from White Papers

FY13 marked the first of the five option years of the program and the initiation of a white paper process for submitting and selecting research tasks to run for 2 to 3 years. Four JX tasks were selected through the white paper process for 2-year durations, focusing on 3 themes that emerged as important in previous years: robust MAV flight in gusts, mapping and exploration by robot teams, and autonomous landing for persistent surveillance. The tasks were as follows:

1) ***Wind/Gust sensing and mitigation on MAST-scale MAVs (Electronics and Mechanics centers and ARL):*** Novel airflow sensors developed by the Microelectronics center were integrated into MAVs and used to estimate gust strength and show the ability to use this information to improve control response to gusts.

2) ***Active cooperative mapping by MAVs (Processing center):*** Methods were developed and demonstrated to establish a common reference frame for multiple robots mapping an environment from different starting points, sharing low data volume information about landmarks detected with 2-D scanning LIDAR sensors.

3) ***Heterogeneous robot teams for mapping 3-D indoor environments (Processing center and ARL):*** A mix of air and ground vehicles were used to do full 3-D mapping of a multi-story building environment, using LIDAR and RGB-D depth camera sensors, where all measurements were communicated to a centralized map.

4) ***Autonomous Rooftop Landing (Integration and Processing center):*** Onboard landing site perception work from previous years was brought together with more sophisticated onboard trajectory planning for jointly optimizing creation of a terrain map and making progress toward a landing site that could observe a prespecified surveillance target.

The Microelectronics center was developing miniature, hair-like sensors to estimate airflow[7] (Fig. 5-3). In Task 1, one of these was integrated into a quadrotor to evaluate the ability to estimate gust velocity and direction as the rotors were operating and the quadrotor was in-flight. Successful measurements were made, but the sensors were fragile; as a result, work in this direction was curtailed pending further maturation of the sensors themselves.

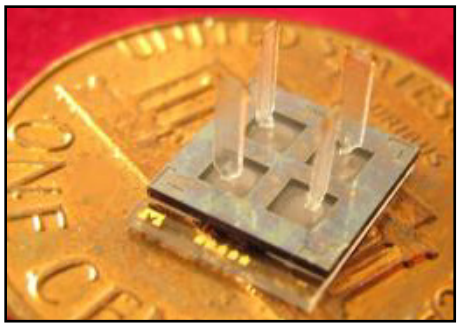


Fig. 5-3 Hair airflow sensor, on a penny to show scale

Tasks 2 and 3 took complementary approaches to cooperative mapping by robot teams. Task 2 restricted mapping to 2-D and studied distributed methods for robots

to establish common reference frames by communicating landmarks from their own map to the other robots, so that each robot could independently estimate its pose relative to the others.[8] Task 3 created full 3-D maps, with a centralized pose graph (one of the robots) for the global map[9]; research included information-based exploration and multi-robot, air-ground mapping experiments (Fig. 5-4, left). This work was generalized in following years to decentralized mapping with algorithms that were shown to scale relatively well to support dozens of robots (Section 5.3.3).

Task 4 extended work from the base years of the program on using SfM to find landing sites on real (outdoor) or surrogate (indoor mock-up) rooftops. A dense SfM pipeline and a trajectory planner that balanced exploration, safety, and progress to the goal were developed and tested in indoor and outdoor experiments[10] (Fig. 5-4, right).

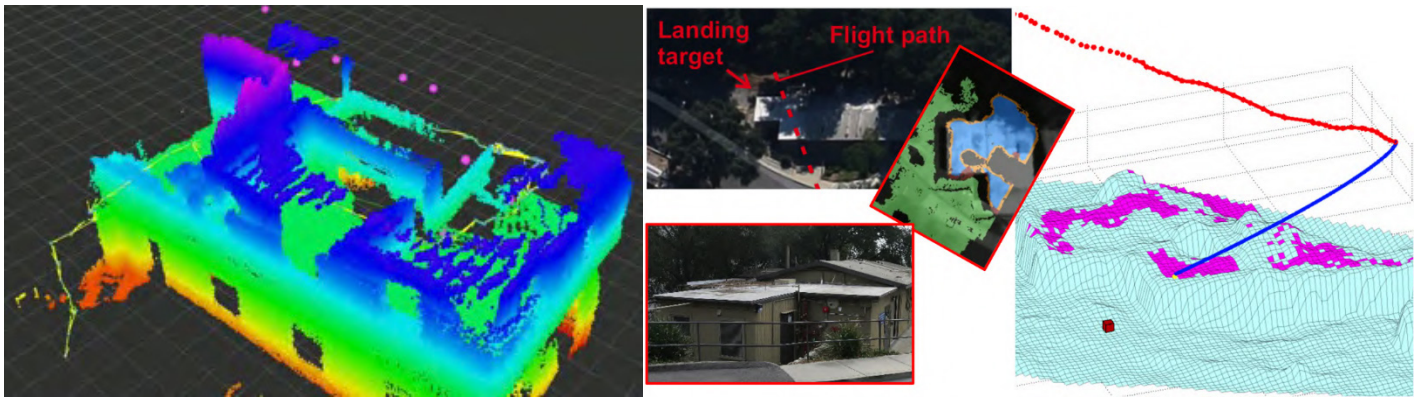


Fig. 5-4 Sample results from FY13–14. (Left) Rendering of a building after being mapped by a team of one air and one ground robot using a centralized map on the ground robot, based on Trevor et al.[9] (Right) Collage of results computed offline from experiment on outdoor landing on a rooftop[10]: (upper) flight path of a data collection over a building; (lower) ground-level photo of the building; (inset) image from a nadir-pointed camera on the quadrotor, with color overlays showing inferred safe (blue) and unsafe (orange) landing sites on the rooftop, and surrounding ground-level terrain (green); and (far right) reconstructed mesh map of the building, manually specified point on the ground to surveil (red dot), actual flight path of the data collection (red line), and trajectory produced by the planner to a safe landing site overlooking the surveillance target (blue line).

5.2.4 Stage 4: FY15–17, Capability Maturation

For the final 3 years of the program, the JX thrust selected a set of topics to focus on to mature and evaluate as key legacies of the program. The motivating scenarios for the program were reexamined to identify technical areas with important capability gaps, but with promising approaches available in the core research program that could be matured and integrated to exercise system-level capabilities in relevant vignettes. With some adjustment after FY15, these tasks focus in the following areas:

1) ***High-speed obstacle avoidance (Processing and Integration center):*** High-speed movement of small robots among obstacles remains a challenge. This task brings together novel perception and planning techniques emerging from the core research thrusts to address this problem.

2) ***Autonomous MAV perch-and-stare (Integration and Processing center):*** This task generalizes work from previous years to improve perception, motion planning, and control for autonomous landing and flight close to structures in more general conditions (e.g., day/night operation and gusts).

3) ***Mapping with heterogeneous robot teams (Processing center and ARL):*** Operating teams of manned and unmanned systems is a major ongoing theme. This task extends prior work to address decentralized SLAM algorithms that scale efficiently to support teams of dozens of robots.

4) ***Airborne sense-and-avoid (SAA) for MAVs (Integration and Processing center):*** All prior work on obstacle avoidance for MAVs addressed static environments. This task addresses onboard perception and planning algorithms to enable MAV flight around other moving aircraft.

5) ***Autonomous navigation on MAST-scale rotorcraft (Integration and Processing center):*** Autonomous navigation capabilities generally have been developed on relatively large testbed aircraft weighing 500 g or more. This task migrates autonomous navigation capabilities into a progression of smaller MAVs.

Table 5-1 summarizes how the state of the art in the areas when these tasks began compares to the current technical vision of the JX thrust. Section 5.3 discusses current work and recent progress on these tasks in greater detail.

Table 5-1 Technical challenges of current JX tasks

State of the art	Thrust area vision
Low-speed obstacle avoidance	High-speed obstacle avoidance
Fragile, daytime autonomous landing with no disturbances	Robust day/night autonomous landing with disturbances
Centralized mapping and exploration by air/ground robot teams of 2 vehicles	Distributed mapping and exploration by air/ground robot teams of 10 or more vehicles
No airborne SAA for MAVs	Short-range airborne SAA for cooperative MAVs
Autonomous MAVs >500 g	Autonomous MAVs <100 g

5.3 Technical Plans and Progress in the FY15–17 Phase

5.3.1 High-Speed Obstacle Avoidance

Many bottlenecks have limited robot velocity with obstacle avoidance:

1) Unreliability of state estimation subsystems in challenging lighting conditions.
2) Low frame rate of perception systems for obstacle detection.
3) Computational complexity of algorithms for world modeling and collision checking.
4) Conservative nature of algorithms for motion planning.

This task directly addresses the last two of these issues; the first two are addressed elsewhere in the MAST program.

In the core research program, we introduced a new approach to world modeling and collision checking that maintains world model representations in image space, instead of the traditional approach of creating and updating voxel-based 3-D world models.[11] We showed that a configuration space (c-space) expansion can be applied to depth maps in image space, and that this allows very efficient collision checking by projecting candidate trajectories into image space and comparing the distance at each pixel along the trajectory to the distance in the depth map (Fig. 5-5). This approach has enabled autonomous flight among obstacle in groves of trees (Fig. 5-6). In these experiments, velocity was limited to 1 m/s due to reliability limitations of visual-inertial state estimation in these difficult lighting conditions and the modest frame rates (e.g., 5 Hz) of software implementations of stereo vision on small processors. Both of these limitations are being alleviated by multiple ongoing developments in this project and elsewhere.

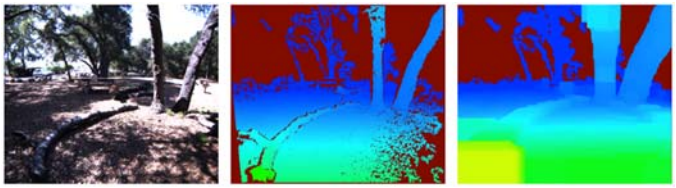


Fig. 5-5 (Left to right) Image of a stereo image pair, depth map from the stereo pair, and c-space expansion of the depth map for efficient collision checking in image space

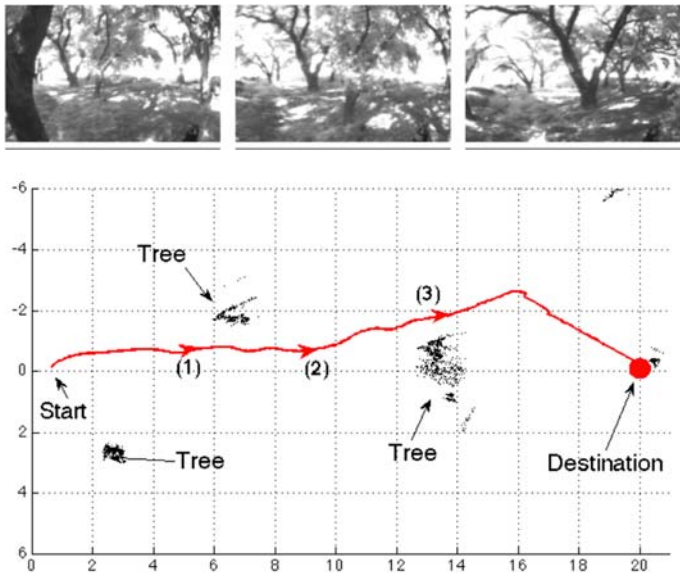
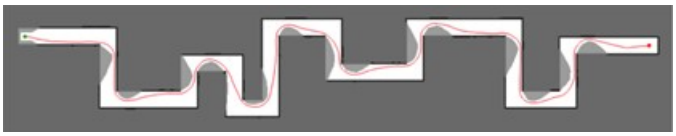
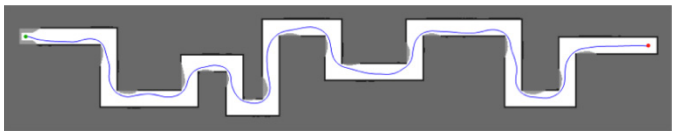


Fig. 5-6 Results of a 20-m experimental flight through a grove of trees, using image space collision checking.[13] (Top) Image from the stereo vision system as three tree obstacles were encountered. (Bottom) Top-down 2-D plot of the trajectory and nearby obstacle pixels from the depth map over the whole run.

Traditional receding horizon motion planners can replace the use of voxel maps with our approach to reduce the computational cost of collision checking for candidate trajectories out to the distance of the planning horizon. Receding horizon planners evaluate candidate trajectories with a cost function that includes a term that estimates the cost to the goal from the end of the planning horizon. We have shown that machine learning can be used to improve estimates of the path cost for the next planning cycle beyond the current horizon.[12] This effectively causes the planner to choose trajectories that take wider turns with higher velocity, which reduces the overall time to reach the goal and reduces the probability of choosing trajectories that require stopping to prevent collisions. We have tested this approach with a small radio-controlled (RC) car that can achieve high velocities safely for rapid development. Experimental results obtained in indoor hallways demonstrate the tendency to take wider turns (Fig. 5-7). Quantitative performance comparisons to a baseline planner without learning depend on the sensor FOV and distance to the planning horizon. Table 5-2 shows comparisons for 2 FOVs with a horizon of 1.5 m.



Baseline planner.



Our planner.

Fig. 5-7 Simulation results comparing trajectories from a baseline planner (top) and our new planner (bottom), showing that the new planner takes wider turns around corners, which reduces the unseen area around the corners (gray regions in top image) and enables shorter time to reach the goal with higher probabilities of success (Table 5-2).

Table 5-2 Experimental results with and without learning

FOV	Planner	Success	Mean Speed (m/s)	Max Speed (m/s)
60°	Baseline	0/10	2.25	4.29
	Learned	**8/10**	**3.08**	**5.35**
90°	Baseline	4/10	3.06	5.30
	Learned	**10/10**	**3.43**	**5.55**

Ongoing work is pushing in several directions is combining image-based collision checking in the front end with the learning-based trajectory evaluation in the back end, as well as investigating three further directions:

1) Extending the image-based depth representation to include temporal fusion in image space to reduce the probability of false alarm obstacles, fill in gaps in depth data that occur in individual frames, and provide a persistent, 360° representation of the nearest obstacle using an "egocylinder" representation.[13,14]

2) Using learning from appearance data in the raw images to improve collision prediction beyond the range of accurate depth perception.[12]

3) Using the depth data at short range to do on-the-fly training of a classifier for real-time segmentation of more distant parts of the image into ground and potential obstacle structures to facilitate motion planning beyond the reliable range of the depth sensor.[15]

5.3.2 Autonomous MAV Perch-and-Stare

MAVs require the ability to land autonomously on a variety of structures for several tactically important reasons. However, often there are difficulties in perceiving the

landing site, maintaining accurate, reliable aircraft position and velocity estimates all the way to touchdown, and coping with external disturbances due to wind and ground effect. Night operation is also extremely desirable, but has not been addressed previously. This task addresses these challenges by extending work from prior years to the following:

1) Adapt trajectories and control policies to permit precise operation near the rooftop surface despite varying wind conditions, surface airflow patterns, and state uncertainty.

2) Operate at night with covert sensors.

3) Perform accurate, reliable, visual-inertial state estimation all the way to touchdown.

Regarding the first goal, our prior work[10] established an informed trajectory generation approach as the state of the art for enabling efficient and robust site detection and landing. However, executing these trajectories in real-world environments is challenging due to uncertainty in sensing, vehicle dynamics, and environmental conditions. Many standard feedback control methodologies used on MAVs (e.g., proportional-derivative feedback control, L_1-adaptive control) do not explicitly account for limits on actuation or system dynamics. Consequently, they cannot mitigate the effects of these sources of uncertainty safely or must yield conservative motion to satisfy large robustness bounds, which conflicts with the objective of landing quickly to conserve power. More advanced techniques, (e.g., nonlinear model predictive control [NMPC]) can mitigate these effects but are impractical for use on a platform with limited computation due to strict SWaP constraints.

We have addressed these limitations by developing computationally efficient methods to generate control policies that enable accurate trajectory tracking. Our nonlinear partial enumeration (NPE) technique uses infrequent optimization to construct a database of control policies over time, allowing it to recover the constraint satisfaction and forward-prediction properties of NMPC while greatly reducing the computational requirements.[16] We have recently addressed wind effects by using the locally weighted projection regression (LWPR) algorithm to extend the approach to be a model-adaptive control technique, so that the control policy database also captures past environmental effects. Most recently, we formulated the Experience-driven Predictive Control (EPC) algorithm, with a faster approach to optimization for adding new control policies to the database.[17] This approach is able to query its controller database with a mean time of 0.77 ms with a variance of 0.75 ms, which reliably enables a 200-Hz control rate on smartphone-

class embedded processors. By comparison, computing a controller via quadratic programming takes 6 times longer with a considerable variance, which does not allow 200-Hz control. Simulations and live experiments show that this approach reduces trajectory tracking error compared to existing methods, while satisfying control constraints (Fig. 5-8).

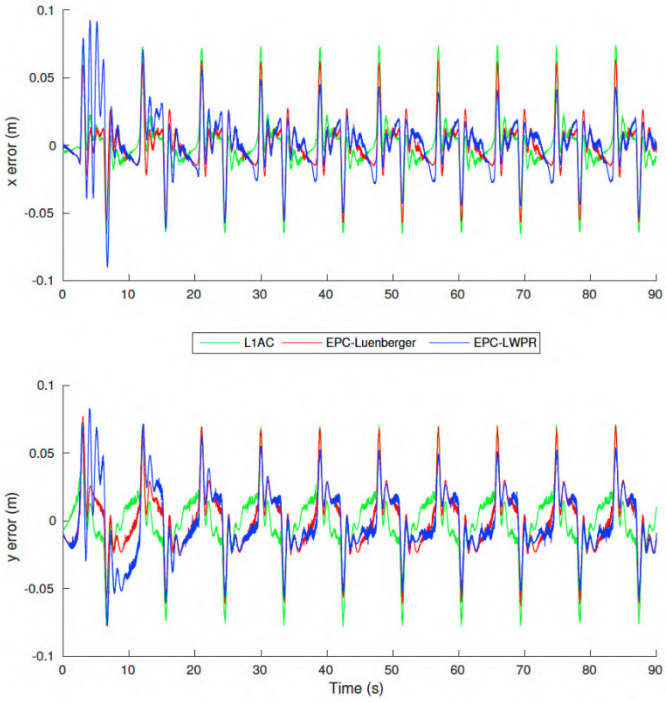


Fig. 5-8 In simulations of a quadrotor repeatedly flying a loop, where part of the loop encounters wind, EPC with LWPR yields improved position tracking error compared to L_1 adaptive control (L1AC) and EPC with a simple state predictor (EPC-Luenberger)

For covert night operation, we developing algorithms to use compact, uncooled thermal IR images for visual-inertial state estimation and 3-D landing site perception via structure from motion (Fig. 5-9). Ongoing work is addressing accurate and reliable position and velocity estimation all the way to touchdown by improvements to visual-inertial state estimation algorithms.

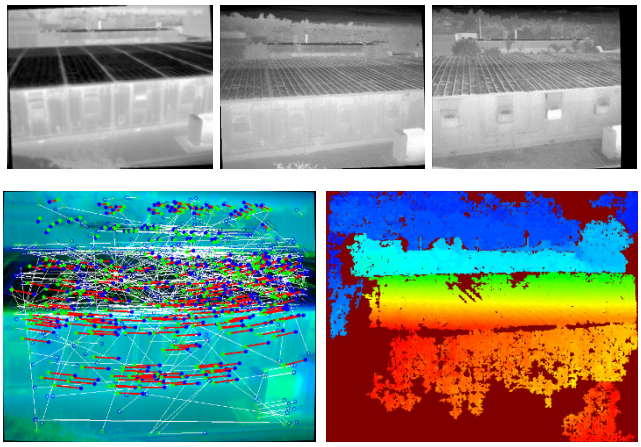


Fig. 5-9 Initial proof-of-concept results for MAV navigation at night with compact, low-cost, uncooled thermal cameras. Thermal image sequences were acquired by manually translating a camera laterally while observing a building. (Top row) One image from sequences obtained at 5 am, 9 am, and 2 pm. (Bottom left) Feature tracking results are good enough to enable processing pairs of images from the sequences as stereo pairs. (Bottom right) False color depth map from dense motion stereo processing (red is closest, blue farthest, maroon indicates no data); with temporal aggregation in an elevation map, such depth maps should enable landing site detection.

5.3.3 Mapping with Heterogeneous Robot Teams

Localizing all members of a team in GPS-denied environments is a critically important capability. Simultaneously mapping the environment both enables and is a valuable byproduct of team localization, hence the importance of algorithms for SLAM. This requires establishing a common reference frame for robots that may start in different locations without knowledge of each other's locations, and communication between robots to share location information.

The backbone of modern SLAM techniques is *pose graph optimization*, which estimates the trajectory of the robots, from which the map can easily be built. The focus of this task is developing algorithms to enable distributed pose graph optimization and reducing the amount of communication required between robots in doing so. We have made two key contributions to this area:

1) Developing algorithms for distributed optimization of the pose graph, which converge to the same answer as a centralized solution that runs on one map server.

2) Developing new methods based on recognizing objects seen by multiple robots to greatly reduce the memory and communication requirements for distributed SLAM.

Figure 5-10 (top left) illustrates the core concept in distributed pose graph optimization. It shows the trajectories of two robots, which occasionally come to

locations (*rendezvous points*, also called *separators* in the pose graph) where they can share information to estimate their relative poses at those points. The crux issues in distributed SLAM are what information to share at those points and how to formulate and solve the joint optimization problem.

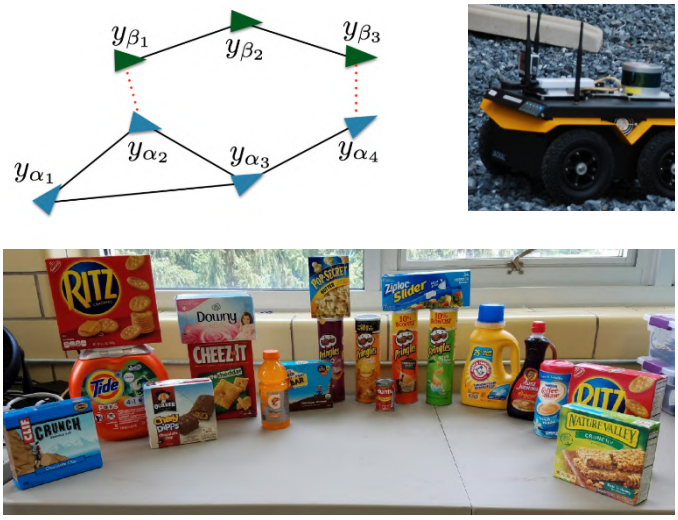


Fig. 5-10 Multi-robot mapping. (Top left) Schematic of pose graph for two robots; dashed red lines indicated where the robots rendezvous and communicate. (Top right) Testbed vehicles for field experiments are Clearpath Jackal robots with a variety of sensors; shown here with a Velodyne LIDAR for the point cloud-based approach. (Bottom) Objects used for object-based approach.

A key observation is that the robots only need to share information about relative poses at the separators; they do not need to share information about their own poses between separators. One class of solutions we have developed shares features from LIDAR-based point clouds obtained by sensors on each robot at separator locations. This involves communicating a data volume that is linear in the number of separators, compared to prior methods that were quadratic in the number of separators.[18] Our most effective optimization technique is an iterative algorithm called the distributed Gauss-Seidel (DGS) method. Simulations have shown good performance with this algorithm with manageable growth in computational cost as the number of robots increases up to several dozen.[19] Analysis, simulation, and field tests with real robots show that the distributed solution converges well to the same answer as a centralized solution (Fig. 5-10b, top right, and Fig. 5-11).

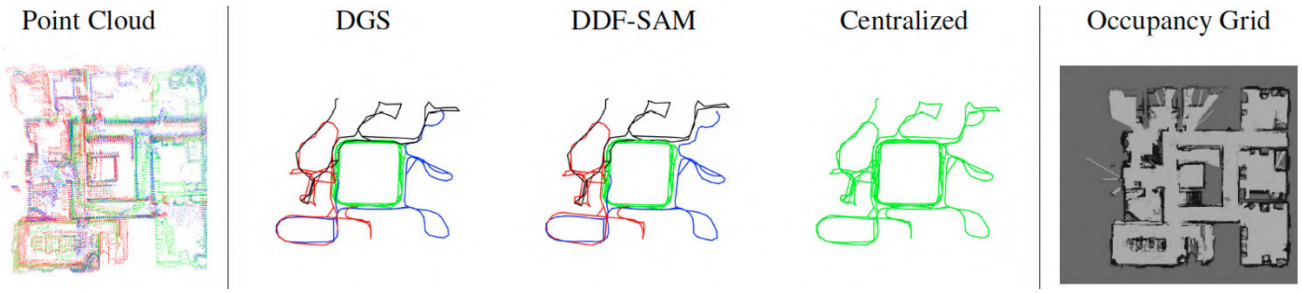


Fig. 5-11 Multi-robot mapping results obtained with four robots indoors at military training facility. (Left) Aggregated point cloud obtained from the DGS trajectory estimate. (Center) Estimated trajectories for DGS, DDF-SAM, and centralized mapping algorithms (robot trajectories shown in different colors for the distributed algorithms). (Right) Overall occupancy grid map obtained from the DGS trajectory estimate. The DGS results are essentially indistinguishable from the centralized approach and the previous state of the art in distributed approaches.

Our second contribution achieves a further reduction in the amount of communication required. If different robots can each recognize the same object at each separator location and estimate their own pose relative to the object, then the only information that must be shared is the pose of each robot relative to the object, rather than the much larger volume of data required to register point clouds obtained by each robot at the separators. We have developed and evaluated such a method, where the 3-D models and appearance of the objects are known to each robot in advance; Fig. 5-10, bottom, shows the objects used so far.[20] Experiments have been conducted with the same robots as shown in Fig. 5-10, using an Asus Xtion RGB-D sensor for object recognition instead of the Velodyne LIDAR. Compared to the method based on registering point clouds at separators, this approach reduces the memory requirements by three orders of magnitude and the communication data volume by six orders of magnitude.

The next step in research in this direction will be to extend the object-based approach to be able to create and recognize object models on-the-fly, rather than to require that all objects and their 3-D models be known in advance. Other work in progress is to extend the optimization algorithm to cope with outliers in the measurement data and apply the entire framework to aerial vehicles, as well as the ground vehicles used to date.

5.3.4 Airborne Sense-and-Avoid for MAVs

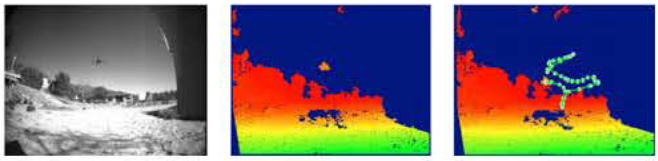
Most research on autonomous navigation for MAVs has addressed operation in static environments or, where multiple MAVs were involved at once, used external motion capture cameras to provide high-quality state information for all aircraft simultaneously. Capabilities for onboard visual-inertial state estimation and vision-based static obstacle avoidance have reached a stage where it is possible to

contemplate aircraft operating in close proximity, using onboard perception to detect and estimate the state of other aircraft.

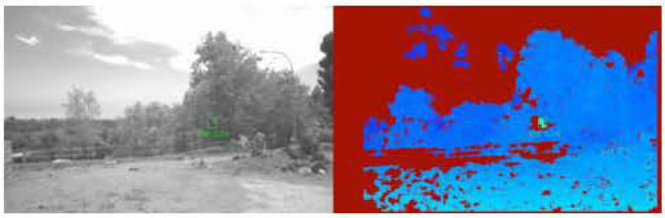
For much larger aircraft, onboard autonomy for avoiding collisions with other aircraft is referred to as the SAA problem. SAA can be considered as "cooperative", where the aircraft involved share sensor data or are aided by ground assets to varying degrees, or "non-cooperative", where each aircraft uses only its own onboard sensors to detect other aircraft and the aircraft do not communicate about their intents. For applications that are relevant to MAST, the non-cooperative case is the most interesting, since it is the stealthiest, but also the most challenging.

Non-cooperative SAA for large aircraft requires long-range detection, using radar or monocular vision methods that work best when the other aircraft is silhouetted against the sky.[21,22] For MAST-scale MAVs, this onboard perception problem had not been tackled prior to our work. MAST-scale radar is for this task does not exist. Monocular vision methods are error-prone when the other aircraft is seen against clutter and do not provide absolute range or velocity of the other aircraft without additional information, such as known size of the other aircraft. For large aircraft SAA, algorithms exist for motion planning given state information about each aircraft.[21] For MAVs, there has been some work on collision avoidance with multiple aircraft in cooperative[23] and non-cooperative scenarios,[24] but only in simulation or with external motion capture systems.

Since onboard real-time stereo vision is available for MAVs, is already in use for static obstacle avoidance (Section 5.3.1), and can provide absolute range and velocity directly within the range limit of the stereo vision system, this is a valuable approach to begin to develop a practical MAV SAA solution. We have extended our real-time stereo vision algorithm to segment small, mid-air objects from the depth map, track them from frame to frame as the host aircraft moves, and estimate the velocity of the other aircraft in an inertial reference frame (Fig. 5-12). This runs onboard a 1.3-kg quadrotor, with a 1.86-GHz dual core Intel Core2Duo processor, in about 120 ms/frame for stereo matching 752 × 480 images, plus 20–30 ms/frame for everything to do with tracking, for a net tracking frame rate of 7 Hz.



Results from a stereo camera pair on a tripod, illustrating tracking a moving quadrotor by segmenting depth maps. Left: one image from the stereo pair showing the quadrotor in the middle of the image. Center: false color depthmap; green is closest, red furthest, and blue is no data. Right: depthmap with overlaid track of moving quadrotor; note detection below the skyline.



Results from a stereo pair on a moving host quadrotor. Left: green box overlaid on onboard image marks where the target quadrotor was detected in front of trees. Right: false color depth map, where light blue is close, dark blue is farther, and maroon represents no data. The small bright path in front of the trees is the quadrotor at closer range than the trees, with a green box overlaid showing that it was detected.



Right: two frames from a data log, showing just the lower right corner of the image for magnification. The host quadrotor was turning to track the target between these frames, causing the change in positon of the background. The green boxes show successful detection in front of the trees. In all images, the green text gives the frame number and the range in meters, which is roughly 5 to 6 m. The maximum detection range was about 15 m with the current camera configuration (20 cm baseline, 80° field of view, 752x480 images).

Fig. 5-12 Stereo-based detection and tracking of moving MAV

To address planning and control, we are building on the QuadCloud aerial swarm framework previously developed under MAST.[24] This includes a high-level planner that provides trajectories for each robot; however, localization errors, external disturbances, or other factors may still lead to collisions. This requires that the trajectories be modified in real time to avoid collisions, while minimizing deviation from the predefined trajectory. Among the many local reactive planning schemes, the reciprocal collision avoidance approach[25] is the most suitable because it requires no cooperation or explicit communication between the aircraft. We have tested this in simulation in preparation for integration with the perception system (Fig. 5-13). We have also developed a centralized multi-robot planning and trajectory

generation method that explicitly accounts for higher-order dynamics, but requires communication and exchange of state estimates between the robots.[26] In experiments, we have demonstrated successful execution of solution trajectories and verified the robustness of the robots' trajectory tracking to unmodeled external disturbances, including the aerodynamic interactions between co-planar neighbors. We have also extended the algorithm to 3 dimensions where quadrotors navigate away from the downwash of their neighbors to improve safety in 3-D workspaces.

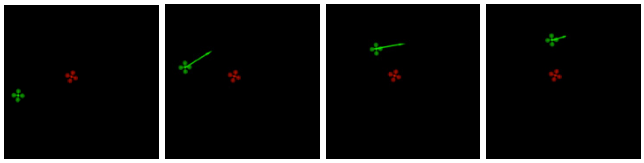


Fig. 5-13 Simulation of reactive collision avoidance algorithm where the red quadrotor is hovering the middle, while the green one is sent a goal on the right and has to maneuver around the red one

5.3.5 Autonomous Navigation on MAST-Scale Aircraft

For much of the MAST program, autonomous navigation capabilities typically were developed on quadrotors of 1 kg more to have sufficient payload capacity to carry the necessary sensors and processors. With several autonomous navigation capabilities now being relatively mature, it is desirable to demonstrate these capabilities with smaller aircraft representing an eventual pocket-sized MAV. The key challenge for this has been the size of the sensors and processors required. Processor miniaturization for the smartphone market has been a major benefit to small, autonomous MAVs, enabling autonomous localization and safe landing using only a camera and inertial sensors in 500-g aircraft several years ago.[27] More recently, major electronics companies have become interested in the drone market, producing tightly integrated sensor/processor boards for drones and extending their software libraries to include highly optimized modules for drone navigation.

After surveying sensor, processor, and rotorcraft airframe options from sources within the MAST program and from the commercial market, we determined that the Snapdragon Flight sensor/processor module designed by Qualcomm offers the best miniaturization and performance option within the remaining timeframe of the MAST program. This includes module a Snapdragon 801 system-on-a-chip processor with an ARM-based, 2.26-GHz quad-core ARM processor for high-level navigation functions, a digital signal processer (DSP) for low-level flight control, an IMU, a barometric altimeter, and nadir and forward-pointed cameras in a board weighing under 20 g. We have implemented algorithms for real-time visual-inertial state estimation and trajectory generation and control on this board (Fig. 5-14, left).

The board has been integrated into quadrotors as small as 250 g; with these, we have demonstrated onboard localization relative to known objects in the environment, followed by onboard planning and execution of aggressive maneuvers through known environments, such as through narrow window (Fig. 5-14, right).[28]

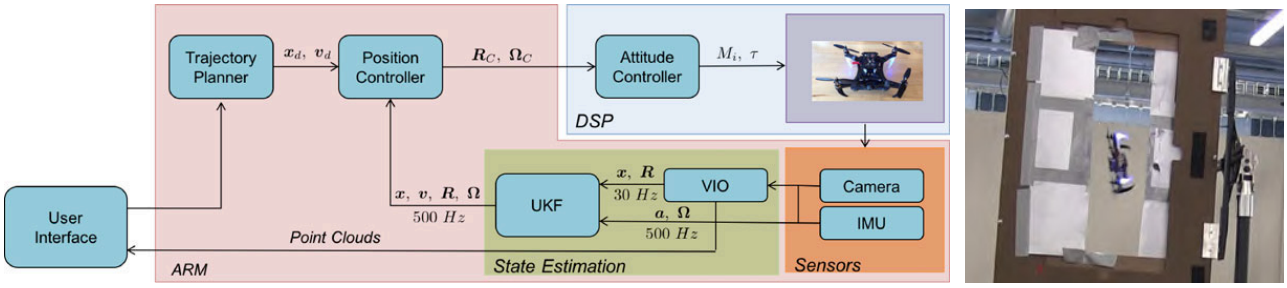


Fig. 5-14 System architecture diagram

Integrating additional capability at this size scale is still limited by the onboard sensors and computing throughput available. Using the nadir-pointed camera on the Snapdragon Flight board, we are integrating our previously developed algorithms (Section 5.3.2) for autonomous landing on building-like structures.[10,27] Before the end of FY17, we expect to demonstrate this capability on a palm-sized quadrotor weighing about 150 g (Fig. 5-15). This board also has a high-resolution, forward-looking camera (3840 × 2160 pixels) and interfaces for a stereo pair of forward-looking cameras, which are available in a smartphone form factor with 640×480 pixels. In principle, these can also be used for obstacle avoidance; this requires more software optimization than we expect to be able to conduct in the remainder of the MAST program.



Fig. 5-15 A 150-g quadrotor for ongoing integration of autonomous landing capability

Further miniaturization and performance enhancement of sensors and processors for autonomous navigation is now progressing rapidly in industry for a variety of commercial applications. Working with multiple industry partners and sponsors,

we continue to integrate more autonomy with such hardware for autonomous navigation of small rotorcraft.

We expect that this will use all of the computing capability of this board, so further advances in autonomy at this scale will require further progress in onboard processing.

5.4 Challenges for the Future

Progress in many areas has been dramatic and is accelerating, fueled by advances in algorithms within the MAST program and rapid advances in compact, high performance sensor, processor, and communications technology for consumer applications, as well as by the surging interest in autonomous mobility for commercial markets. The limited endurance of small robotic platforms remains a bottleneck due to limitations of onboard energy storage capacity. This is being alleviated to some degree by novel concepts of operations that are enabled by more onboard autonomy and will be further alleviated by advances in power and energy technology driven by large commercial markets.

Research in the JX thrust is contributing to key foundations for onboard visual-inertial state estimation; autonomous navigation among static and moving obstacles; distributed SLAM algorithms for large, heterogeneous robot teams, autonomous landing for MAV perch-and-stare missions; and other capabilities. More research is required in many directions, including improving the robustness of perception systems in extreme lighting conditions, increasing robot speed to match the operational tempo of Army missions, and embedding more semantic understanding of objects around robots into autonomous systems. On-the-fly learning will be an important component of several of these directions.

5.5 JX Task Summary

CENTER LEAD: Larry Matthies -Jet Propulsion Laboratory			
ARL TECH LEAD: David Baran			
Project #	**Project Title**	**PI**	**Org**
JX-16-1	Aggressive flight using approximate inverse depth models	Roy / Matthies	MIT / JPL
JX-16-2	Autonomous MAV perch-and-stare toward persistence surveillance	Matthies / Michael	JPL / CMU
JX-16-3	Heterogeneous Robot Teams for Mapping of Environments	Christensen / Kumar	GaTech / UPenn
JX-16-4	Cooperative Aerial Reconnaissance of Moving Objects	Kennedy / Kumar	JPL / UPenn
JX-16-5	Integrating Autonomous Navigation on MAST-scale Rotorcraft	Brockers/ Kumar	JPL / UPenn

5.6 References

1. T. Watteyne, F. Chraim, N. Sarmicanic, C. Jian, and K. Pister, "Video transmission over a standards-based wireless multi-hop sensor network," E-Letters of the IEEE ComSoc Multimedia Technical Committee, Special Issue on "Multimedia Over Embedded Systems", May 2010.

2. Kamal Sarabandi, Jungsuek Oh, Leland Pierce, Kunigal Shivakumar, Shivalingappa Lingaiah, "Lightweight, Conformal Antennas for Robotic Flapping Flyers," IEEE Antennas and Propagation Magazine, Volume: 56, Issue: 6, Dec. 2014.

3. Jungsuek Oh, Kyusang Lee, Tyler Hughes, Stephen Forrest, and Kamal Sarabandi, "Flexible Antenna Integrated With an Epitaxial Lift-Off Solar Cell Array for Flapping-Wing Robots," IEEE Transactions on Antennas and Propagation, Vol. 62, No. 8, August 2014.

4. Oh, J. (2012). *Advanced Antenna Miniaturization Techniques for Low-Power, Broadband and Diversity Wireless Systems* (Doctoral dissertation, University of Michigan).

5. S. Zarovy, M. Costello, and A. Mehta, "Experimental method for studying gust effects on micro rotorcraft," Journal of Aerospace Engineering, 227(4), 2013.

6. R. Gardner and J. S. Humbert, "Comparative framework for maneuverability and gust tolerance of micro helicopters," Journal of Aircraft, 51(5), 2014.

7. M. Sadeghi, R. L. Peterson, and K. Najafi, “A 2-D directional air flow sensor array made using stereolithography and MEMS micro-hydraulic structures,” 17th International Conference on Transducers, 2013.

8. V. Indelman, E. Nelson, N. Michael, and F. Dellaert, “Multi-robot pose graph localization and data association from unknown relative poses via expectation maximization,” Int’l Conf on Robotics and Automation, 2014.

9. A. J. Trevor, J. G. Rogers, and H. I Christensen, “Omnimapper: a modular multimodal mapping framework,” Int’l Conf on Robotics and Automation, 2014.

10. V. R. Desaraju, N. Michael, M. Humenberger, R. Brockers, S. Weiss, J. Nash, and L. Matthies, “Vision-based landing site evaluation and informed optimal trajectory generation toward autonomous rooftop landing,” Autonomous Robots, vol 39, 2015.

11. L. Matthies, R. Brockers, Y. Kuwata, and S. Weiss. “Stereo vision-based obstacle avoidance for micro air vehicles using disparity space,” Int’l Conf on Robotics and Automation, 2014.

12. C. Richter and N. Roy, “Learning to plan for visibility in navigation of unknown environments,” International Symposium of Experimental Robotics, 2016.

13. R. Brockers, A. Fragoso, B. Rothrock, C. Lee, and L. Matthies, “Vision-based obstacle avoidance for micro air vehicles using an egocylindrical depth map,” International Symposium of Experimental Robotics, 2016.

14. C. Cigla, R. Brockers, and L. Matthies, “Gaussian mixture models for temporal depth fusion,” Winter Conference on Applications of Computer Vision, 2017.

15. M. Bajracharya, A. Howard, L. Matthies, B. Tang, and M. Turmon, “Autonomous off-road navigation with end-to-end learning for the LAGR program,” J. Field Robotics 26(1): 2009.

16. V. Desaraju and N. Michael, “Fast nonlinear model predictive control via partial enumeration,” Int’l Conf on Robotics and Automation, 2016.

17. V. Desaraju and N. Michael, “Experience-driven predictive control,” Robot Learning and Planning Workshop, held at the Robotics Science and Systems Conference, 2016.

18. A. Cunningham, V. Indelman, and F. Dellaert, "DDF-SAM 2.0: consistent distributed smoothing and mapping," Int'l Conf on Robotics and Automation, 2013.

19. S. Choudhary, L. Carlone, C. Nieto, J. Rogers, H. I. Christensen, and F. Dellaert, "Distributed trajectory estimation with privacy and communication constraints: a two-stage distributed Gauss-Seidel approach," Int'l Conf on Robotics and Automation, 2016.

20. S. Choudhary, L. Carlone, C. Nieto, J. Rogers, Z. Liu, H. I. Christensen, and F. Dellaert, "Distributed trajectory estimation with privacy and Multi Robot object-based SLAM," International Symposium on Experimental Robotics, 2016.

21. X. Yu and Y. Zhang, "Sense and avoid technologies with applications to unmanned aircraft systems: review and prospects," Progress in Aerospace Sciences, vol. 74, 2015.

22. D. Dey, C. Geyer, S. Singh, M. Digioia, "A cascaded method to detect aircraft in video imagery," International Journal of Robotics Research 30(12), 2011.

23. K. Mohta, M. Turpin, A. Kushleyev, D. Mellinger, N. Michael, and V. Kumar, "QuadCloud: a rapid response force with quadrotor teams," Int'l Symp on Experimental Robotics, 2014.

24. J. Alonso-Mora, T. Naegeli, R. Siegwart, and P. Beardsley, "Collision avoidance for aerial vehicles in multi-agent scenarios," Autonomous Robots, 39(1), 2015.

25. D. Bareiss and J. van den Berg, "Reciprocal collision avoidance for robots with linear dynamics using LQR-obstacles," Int'l Conf on Robotics and Automation, 2013.

26. S. Tang and V. Kumar, "A complete algorithm for generating safe trajectories for multi-robot teams," Int'l Symp on Robotics Research, 2015.

27. R. Brockers, M. Humenberger, S. Weiss, and L. Matthies, "Toward autonomous navigation of miniature UAV," IEEE Embedded Vision Workshop, held with CVPR, 2014.

28. G. Loianno, C. Brunner, G. McGrath, and V. Kumar, "Estimation, control and planning for aggressive flight with a small quadrotor with a single camera and IMU," Int'l Conf on Robotics and Automation, 2017.

Autonomous MAV Perch-and-Stare Towards Persistence Surveillance

JX-16-17-2

Vishnu R. Desaraju, Alexander Spitzer, Nathan Michael
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
{rajeswar, spitzer, nmichael}@cmu.edu

Shreyansh Daftry and Larry Matthies
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109
{Shreyansh.Daftry, lhm}@jpl.nasa.gov

***Abstract*—Micro air vehicles require the ability to land autonomously on a variety of structures to enable persistent perch-and-stare missions and other functions. In previous years, we demonstrated autonomous landing in indoor and outdoor experiments, using monocular structure-from-motion with visible spectrum cameras for depth perception to find the landing site and informed trajectory planning to balance landing site detection and rapid landing. Here, we report progress on experience-driven predictive control algorithms that efficiently adapt to the external disturbances that are experienced close to landing and on steps to generalize the perception system to use a thermal infrared camera to enable night operation.**

I. Introduction

A. Background and Technical Challenges

Micro air vehicles (MAV) require the ability to land autonomously on a variety of structures for several reasons, including to enable persistent perch-and-stare missions, to conserve energy during lulls in mission timelines, to recover the vehicle at the end of missions, and to recharge onboard batteries from external power sources. Potential landing sites include a wide variety of rooftop structures (e.g. horizontal versus sloped, smooth versus rough), walls, tree branches, light posts, ground vehicles of various kinds, and other objects. In this task, our main focus is persistent perch-and-stare scenarios involving landing on rooftop-like structures in realistic environmental conditions, including day and night operation with disturbances from wind and the vehicle's own propeller downwash. We envision scenarios where the MAV must initially perceive the landing site from relatively far away (e.g. 10s of meters) and can refine its perception of the site as it approaches. We include scenarios where the vehicle may take up its observing position at night for stealth. Thus, the sensor suite must not only have minimal size, weight, and power, it must also enable rooftop landing site detection from both short and long range, day and night.

In prior work under Joint Experiment 4 (FY13-14), we developed a first generation of a rooftop landing system [1], [2] For perception, this used a single, downward-looking, visible spectrum camera with dense motion stereo algorithms for rooftop detection and landing site safely analysis. The camera was also used for feature tracking in vision-aided inertial state estimation. A basic model of the confidence in the safety of landing sites was embedded a probabilistic landing site detection and selection strategy. We developed an information-based trajectory generation algorithm that balanced objectives to explore for safe rooftop landing sites with reaching a vantage point overlooking a specified surveillance target. The approach captures rooftop surface uncertainty and landing confidence, integrating the vision-based site selection and online planning algorithms into a cohesive framework. Outdoor data sets representing a full mission scenario were processed successfully off-line. The software was fully integrated in a real-time, onboard system [3] and the vision and planning components were tested separately in live operation in indoor environments. An end-to-end live experiment was completed at the end of FY14 in a small-scale, indoor set-up landing on a box as a surrogate rooftop.

Follow-on work in FY15 extended this in several ways. Work on perception addressed outdoor operation, acquiring additional data sets flying at multiple heights and with multiple approach trajectories over three buildings at JPL; high resolution lidar scans of the rooftops were acquired as ground truth shape. Testing identified two key limitations of the FY14 algorithms: (1) the sequence of point clouds generated by structure from motion from successive pairs of images showed registration errors that were not evident in indoor testing, and (2) the visual inertial state estimation algorithm using a nadir-pointed camera was not adequately reliable very close to surfaces, when there sometimes are too few visual features to track. Causes of the first issue were diagnosed and solutions were obtained. The second issue is understood and improvements are in development.

FY15 work on planning and control studied the generalizability and robustness of the landing site selection and information-based trajectory generation methodology to different rooftop scenarios and degraded state information. The approach was found to be effective in experiments with multiple data sets. To address the high computational complexity of optimization-based trajectory generation, we leveraged our earlier work with the Closed-Loop RRT algorithm (CL-RRT),

which also facilitates obstacle avoidance. We added a local occupancy grid representation of the rooftop environment to the Gaussian process (GP)-based landing site preference modeling system, and conducted additional indoor flight tests that showed the system to be effective for landing at sites that are safe, overlook a desired surveillance target, and avoid obstacles on the surrogate rooftop. The whole system was implemented on a mobile-class ODROID computer. Most recently, we have started investigating approaches that combine the real-time performance of the CL-RRT planner and the optimality of the spline trajectory generation approach. This led to a new nonlinear model predictive control (NMPC) technique, called Nonlinear Partial Enumeration (NPE), which achieves significant reductions in runtime.

FY16 development resulted in an Experience-driven Controller (EPC) that further reduces the computational requirements of Model Predictive Control through the use of reusable controller databases. A robust extension to EPC is capable of mitigating state uncertainty for the purposes of constraint satisfaction. Explicitly defined operational constraints allow for guarantees in vehicle performance in the presence of obstacles and actuation limits. Furthermore, we combined EPC with a CL-RRT planner in an offline controller database computation phase to enable high-rate control without online optimization on severely computationally constrained systems. FY16 also saw the validation of Locally Weighted Projection Regression (LWPR) and Incremental Sparse Spectrum Gaussian Process Regression (ISSGPR) as efficient algorithms for the estimation of aerodynamic disturbances.

While substantial progress has been achieved, several key technical challenges remain. Work to the end of FY16 only addressed daytime operation, using a standard daylight camera, for scenes with sufficient texture to allow depth estimation by passive triangulation. Enabling night operation would significantly enhance the tactical value of such a system, but requires different sensors, which require additional algorithm development. Structure from motion algorithms must work reliably in different scenarios, requiring robustness to a wide range of camera motions. Self-contained, visual-inertial state estimation algorithms must be extremely reliable at all times up to touchdown. While the planner and controller have been designed to accommodate and adapt to sensor and actuator uncertainty as well as perceived landing site safety, further work is needed to adapt to disturbances stemming from airflow interacting with the surface and to span the entire operational scenario from relatively high-altitude reconnaissance for landing sites to covert flight very close to the rooftop to land at the ultimate destination.

B. Research Objectives

This task addresses all of the above challenges with the goal of enabling a quadrotor MAV to autonomously:

1) Approach a rooftop and identify and land on a site that ensures visibility of a user-provided interest point using a self-contained, onboard perception system during day or night operation. Since we have previously addressed day operation, the main focus of FY17 work was on night operation with a COTS, low-cost, uncooled thermal infrared camera.
2) Perform reliable visual-inertial state estimation and structure from motion at all altitudes, all the way to touchdown and under a wide range of camera orientations.
3) Adapt trajectories and control policies to permit precise operation near the rooftop surface despite varying wind conditions, surface airflow patterns, and state uncertainty.

II. Related Work

Prior work on autonomous landing has focused on two different scenarios: (1) landing in open terrain among obstacles and (2) landing on ships. Both areas have used much larger aircraft with much larger sensor and processor suites and, in many cases, much more computationally intensive algorithms. The problem of landing a micro air vehicle on rooftops or other elevated structures for surveillance purposes brings different challenges than these other scenarios, from the perspective of both perception algorithm design and miniaturization requirements.

There has been some prior work on perching on vertical surfaces, but this work was focused almost exclusively on trajectory design and control aspects of the problem in benign environments (e.g. no wind, no clutter) with well-known surfaces to land on; there was no attention in that work on perceiving the perch site or coping with wind or viewpoint constraints. To our knowledge, the only prior work that addresses the full problem of landing micro air vehicles on real, elevated surfaces is the work we have done in this consortium.

This task draws on work from previous years in task JX-14-4 and JX-15-3 on autonomous landing. It also collaborates with and draw new results from a supplemental task that includes a collaboration between JPL, CMU, and ARL in task SUPP-15-16-SPP on comparative evaluation of approaches to tightly coupled filters for visual-inertial state estimation.

III. Approach

A. Perception for State-Estimation and Landing Site Detection

Night operation requires a thermal infrared (TIR) camera. COTS, uncooled, TIR cameras based on microbolometer focal planes are now small enough to be practical payloads on MAVs. These cameras have several characteristics that may be challenging for vision algorithms. Nevertheless, our work to date suggests that visual-inertial navigation and landing site detection is possible with these cameras; demonstrating this conclusively for night operation would have tremendous tactical value. Advanced TIR focal plane technology that has been under development for several years in the MAST Microelectronics Center has potential to enable even smaller cameras with better imaging performance.

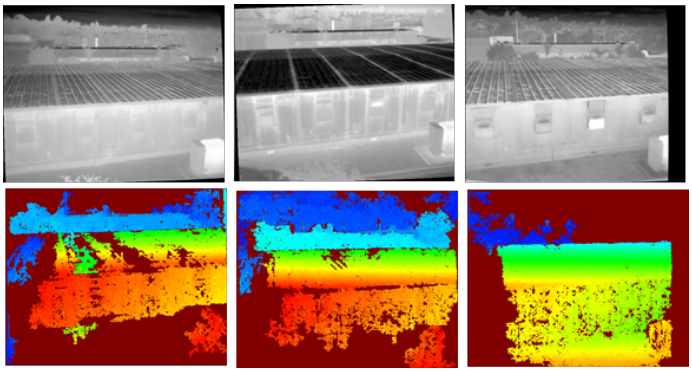


Fig. 1: Stereo reconstruction results for thermal images at different times of the day. (a) Pre-Dawn - 5am, (b) Mid-morning - 9am, (c) Mid-afternoon - 2 pm. Stereo was computed using the Semi-Global Matching (SGM) stereo algorithm.

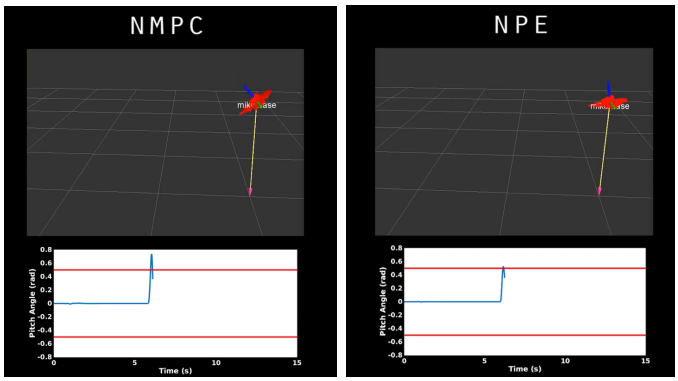


Fig. 2: (a) NMPC often returns suboptimal solutions due to computational constraints, in this case, leading to a large violation of the pitch angle constraint. (b) NPE does not exhibit this failure mode due to its reduced computation time and database of controllers that inherently capture the constraints.

Microbolometers have a time constant that is analogous to the exposure time of photon detecting focal planes (e.g. visible spectrum imagers). Time constants are currently around 10 msec, which is far longer than the exposure times possible with CMOS imagers in sunlight. Under rapid camera motion, this could cause image blur. Microbolometers also use rolling shutter image readout architecture, instead of the global shutter that is almost universally used in MAV vision systems currently. This would lead to geometric image distortion under rapid camera motion. Empirically, in our experience the signal to noise ratio (SNR) of microbolometers often is not as good as good visible spectrum imagery.

Despite all this, our work to date shows promise for doing visual-inertial navigation and landing site detection with TIR cameras. To assess whether stereo reconstruction would be feasible, we took thermal images of a building from a viewpoint that included both the roof and side of the building. This was done for three different times of day and night. Two images with a 1.5m baseline were taken each time, and these were processed within the current pipeline to find stereo reconstructions: features were matched between the images, the relative camera motion was solved for, the images were rectified, and stereo depth images were produced. Figure 1 shows depth maps produced using the Semi-Global Matching stereo algorithm [4]. Adequate results were achieved for both the ground and the building.

Given these results, we expect that visual-inertial navigation and monocular structure from motion for landing site detection will be possible at low flight speed with minimal modification of our algorithms. As flight speed increases, it may be necessary to compensate for rolling shutter artifacts. Visual-inertial navigation (VINS) with visible spectrum rolling shutter cameras in cell phones and tablet computers has been studied in depth, and effective algorithms have been demonstrated with modifications of tightly-coupled VINS filters [5], [6]. Landing site detection requires structure from motion algorithms that provide fairly dense 3-D reconstruction, as illustrated in Figure 1. As flight speed increases, it may be necessary to account for rolling shutter distortion as part of image rectification; techniques for doing so have been developed [7].

The failures of VINS algorithms very close to the surface stem from two main issues. First, depending on the scene, the number and distribution of visual features available for tracking may get worse very close to the surface. Second, this issue can be exacerbated by aiming the camera straight down, which is how many of our data sets have been acquired to date because that works well with dense motion stereo for landing site detection at relatively high altitude. The reliability problem may be alleviated significantly by aiming the camera more forward, allowing observation of features in a larger portion of the scene. The difficulty with having few features or a poor spatial distribution of features is particularly problematic for loosely-coupled VINS filters that require sufficient features to estimate 5 degrees of freedom of camera motion by the vision system alone. Tightly-coupled filters use feature tracking measurements to update the state vector without estimating the full camera motion just from visual features; thus, these are inherently more robust to variations in the feature distribution. Much of the solution to this problem was developed in FY16 under task MAST-FY15-SUPP-WP, which is a collaboration between JPL, CMU, and ARL to develop and compare alternate state estimation methods. Experience in FY17 showed that performance with the MSCKF tightly-coupled filter was a good compromise in cost/performance, with adequate speed on Intel NUC-class embedded processors. However, this algorithm is still somewhat slow (5 Hz) on the smaller Odroid-class embedded processor. Further performance optimization is underway, and we expect to achieve 15+ Hz on Odroid-class processors.

B. Informed online trajectory generation and experience-based control

Our work to date has established an informed trajectory generation approach as the state of the art for enabling efficient and robust site detection and landing. However, executing

these trajectories in real-world environments presents several challenges due to uncertainty in sensing, vehicle dynamics, and environmental conditions. Many standard feedback control methodologies used on MAVs (e.g., proportional-derivative feedback control, $\mathcal{L}_1$-adaptive control), do not explicitly account for limits on actuation or system dynamics. Consequently, they are unable to mitigate the effects of these sources of uncertainty safely or must yield conservative motion to satisfy large robustness bounds, which conflicts with the objective of landing quickly to conserve power. More advanced techniques, (e.g., nonlinear model predictive control, NMPC) can mitigate these effects but are impractical for use on a platform with limited computation due to strict size, weight, and power (SWaP) constraints.

Our recently-developed nonlinear partial enumeration (NPE) technique uses infrequent optimization to construct a database of control policies over time, allowing it to recover the constraint satisfaction and forward-prediction properties of NMPC while greatly reducing the computational requirements, making this approach viable for SWaP constrained platforms. Figure 2 illustrates the MAV starting an aggressive translational motion, resulting in a large commanded pitch angle. With the controller running at 200 Hz to stabilize the MAVs fast attitude dynamics, NMPC often fails to find an optimal solution within 5 ms (mean solution time: 3.8 ms, standard deviation: 0.9 ms), and as a result may violate constraints. However, NPE is able to leverage its database to compute commands with far less computation (mean solution time of 0.075 ms, standard deviation: 0.186 ms) that satisfy the constraints.

To leverage the computational advantages of NPE in a real-world rooftop landing scenario, it must be extended to mitigate the effects of uncertainty. While a MAV operating in any outdoor environment will experience ambient disturbances (e.g., wind), landing inherently requires operation in close proximity to the rooftop, which will induce additional severe disturbances due to the interaction between the MAVs thrust and rooftop geometry. Operation near rooftop surfaces also involves severe variations in environmental effects on the perception system. In addition, the sensors for night operation will yield state estimates with different uncertainty distributions compared to those used for daytime operation.

Given that we wish to enable accurate trajectory tracking, we seek to quantify the effects of these sources of uncertainty on vehicle motion, rather than estimating the actual disturbances. Since NPE, like NMPC, is a model-based control approach, we can leverage online model estimation techniques, such as locally-weighted projection regression (LWPR) [8] and incremental sparse spectrum Gaussian process regression (ISSGPR) [9] to extend the approach to be a model-adaptive control technique. With this formulation, changes in the model capture the varying effects of these sources of uncertainty and lead to a modified control policy that is better suited to controlling the systems current dynamics. Specifically, we aim to generalize the notion of the control policy database in NPE to a database of past experiences, i.e., the states, dynamic models, and sensing models that informed previous control actions. As these models are learned and the corresponding experiences are stored, we can efficiently select control actions that track the landing trajectory accurately, even in the presence of uncertainty in the MAVs motion and sensing. Moreover, these online learned models can be leveraged at the trajectory planning level as well, yielding trajectories that also adapt to the uncertainty in the environment.

C. Integration, Evaluation, and Demonstration

We built upon the existing integrated capability foundation developed through prior MAST tasks including a common simulation and experimentation software and hardware framework to enable rapid prototyping and integration of algorithmic approaches jointly developed through this collaborative task. The first year integration objectives focused on finalizing the end-to-end online, real-time, adaptive planning and control framework and perception system, culminating in experimental evaluation in a controlled environment. We extended these results to consider experimental evaluation of the integrated system when landing on representative outdoor rooftop surfaces under various environmental conditions.

In FY17, we developed capstone demonstration scenarios that highlighted the autonomous landing capabilities developed through the MAST program. The demonstrated capabilities included robust online trajectory generation and model adaptation toward precise approach and near-surface flight as well as day or night vision-based state estimation and landing site evaluation. As discussed below, night operation capabilities were tested onboard with very slow, manually controlled flight, due to the need for further work to address frame-to-frame image gain variability. For the capstone demonstration, we emphasized experimental evaluation of the capabilities and algorithmic extensions that ensure concurrent, real-time performance on SWaP constrained systems.

IV. Results

A. Metrics for Performance Evaluation:

1) Evaluate accuracy, completeness, and reliability of visual structure from motion for perceiving 3-D structure of rooftop and discriminating it from surrounding terrain: Due to our variable baseline selection algorithm, which chooses a baseline that is linear with altitude, elevation map errors are also linear with altitude, instead of the quadratic behavior produced with fixed baselines. With good state estimation, elevation errors are on the order of 1% of altitude. Completeness of depth maps depends on the scene, the algorithm, the camera (i.e. visible or thermal), and the ambient conditions; depth map data density on the order of 90% was observed with well-textured scenes. Reliability here is the percentage of runs without catastrophic failure of structure from motion; this is expected to be $> 80\%$.

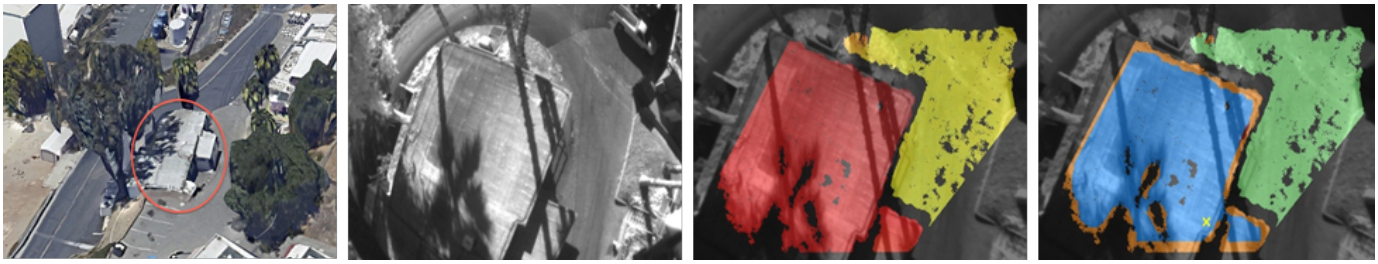


Fig. 3: Real-world landing site detection and autonomous landing using a MAV on an unoccupied building at JPL. 3D maps of the area were generated using the dense motion stereo algorithm, and color-coded based on landing site affordance: (Green) Too low to the ground, (Orange or Red) Unsafe for landing, and (Blue) Safe landing locations on the rooftop.

2) Evaluate the repeatability of the landing site detection, landing trajectory generation, and state estimation components: State estimation was evaluated in terms of error variances across repeated descents in similar scenarios. For the sake of integration and capstone demonstration of outdoor, autonomous landing, trajectory generation was reduced to linear flight segments designed to pass horizontally over the structure to be landed on and to descend vertically to the chosen safe landing point. Repeatability of identifying and landing at a feasible site was characterized by failure rates ($< 20\%$) and spatial distribution of designated landing sites over repeated trials (selected sites across multiple trials of a given scenario lie within a 1.0 m diameter region).

3) Compare trajectory tracking accuracy over multiple trials of the experience-based controller and traditional proportional-derivative and linear MPC with and without varying illumination and aerodynamic disturbances: Accuracy is defined with respect to metric error between the reference and actual performance. We anticipate position tracking accuracy will be within 0.1 m after accounting for state estimation error. The evaluation also included an analysis of the expected computational load for the respective control strategies and energy expenditure toward showing that the proposed approach yields more accurate tracking performance with lower computational load and energy expenditure. The controller should be able to maintain an update rate of 100 Hz while using less than 10% of the available computational power, on average.

4) Assess the computational load due to each perception, planning, and control software module to validate their viability for online processing onboard a MAV: The evaluation included both the average and peak computation time across a range of platforms, including desktop, mobile, and low-power architectures, to assess the scalability of the integrated system to different SWaP requirements. We anticipate the typical computational load will be less than 80% of the available compute cycles, thereby providing a safety margin for fluctuations due to varying operating conditions.

B. Research Accomplishments

1) Landing Site Detection: We evaluated our onboard, monocular landing site detection system by landing on a real building outdoors as shown in Figure 3(a)-(d). The building is located on the JPL campus and is a small, unoccupied building with a flat roof. The quadrotor was autonomously flown above the building, during which time the perception pipeline used motion stereo to generate a 3D map of the area. This map is shown in the Figure 3(d) , where the colors indicate positions too low to the ground (green), unsafe for landing (orange and red), and safe landing locations on the rooftop (blue).

The dense motion stereo algorithm however is constrained to work well for motions parallel to the ground. In order to alleviate this limitation, and extend our SfM capability to tolerate a wide range of motions, we also conducted initial experiments with a direct SLAM based structure from motion algorithm [10], which uses a Kalman filter-style algorithm to incrementally update image-based depth maps from measurements of optical flow between successive images. The experiments were promising, though so far only qualitative. There are trade-offs between wide-baseline, dense motion stereo and incremental depth map estimation from short-baseline optical flow that require further examination. Future work will address quantitative performance analysis of the two SfM algorithms on different environments and landing scenarios. Additionally, to facilitate parallel progress on monocular structure from motion and landing site perception at high altitude from planning work addressing navigation at low altitude, a quadrotor at CMU was equipped with a stereo camera pair for depth estimation at low altitude.

For the capstone demonstration, a sheet of plywood was covered with asphalt roofing shingles and set on a stand roughly 1 meter above the ground to represent a realistic-looking rooftop landing target. Up to six consecutive landing trials were executed successfully with this target. Failures were generally due to state estimation and control idiosyncracies on the relatively low performance Odroid processor, particularly as wind increased.

2) Night Operations with Thermal Infrared:: We integrated a FLIR System's A65 thermal infrared camera on a Asctec Pelican quad-copter, as shown in Figure 4a, and evaluated the onboard performance of visual-inertial state-estimation and landing-site estimation in a natural outdoor environment at different times of the day and night. The results of our analysis showed that, on average, the SURF detector [11] achieved the highest repeatability scores when SNR was low. However FAST features offer the best tradeoff for performance vs real-

time performance. For tracking, we used a sparse iterative version of the Lucas-Kanade (LK) algorithm [12]. Figure 4 shows results of successful feature detection and tracking over a sequence of frames. Further, we also performed a qualitative evaluation for the performance of the landing site estimation pipeline on the TIR modality. Initial experiments verified our belief that landing site estimation is possible at low flight speed with no modification to our algorithm. At higher speeds, state estimation experienced difficulty. We believe this was due to image gain variation as hot objects moved in and out of the image, which was not adequately handed by the LK tracking algorithm implementation that assumed constant brightness. This issue remains to be addressed in the future.

3) Close-range Perception with Active Range Sensors: We evaluated the possibility of using active sensors that rely on a near-infrared emitter collocated with the sensor, such as the Intel RealSense. This allows the MAV to perceive surfaces within about 4 meters distance in low light or darkness, whether the MAV is moving or stationary; this provides greater reliability and operational flexibility than with using just passive SfM with visible or thermal imagery. We studied the performance of the Realsense sensor and derived a range error model for it. While the range error is quadratic in true range for a passive stereo, it was found to be cubic in range for the RealSense:

$$\sigma_z = (\frac{1}{fb})\frac{\sigma_d}{Z_0}Z^3 \tag{1}$$

where f is the focal length, b is the baseline, Z is the range and Z_0 is the control distance for the camera. The sensor model was experimentally validated using data from the RealSense, in lighted and dark conditions, as shown in Figure 5. Furthermore, the sensor was integrated onboard an Asctec Hummingbird quadrotor and end-to-end autonomous landing experiments were performed to show the improved perception capability close-to-ground using the active sensor.

4) Robust Experience-driven Predictive Control: We formulated a robust extension to Experience-driven Predictive Control (Robust EPC) [13] that enables probabilistic state constraint satisfaction in the presence of state uncertainty. As the environment around the vehicle changes, sensing quality may degrade and result in increased vehicle state estimation and localization error. To mitigate the increased uncertainty, Robust EPC intelligently adjusts controller conservatism to ensure that mission critical constraints (e.g. avoiding obstacles) are satisfied with a desired probability. This guarantee can be made without sacrificing performance, as Robust EPC uses a Gaussian belief propagation model to predict state uncertainty over the optimization horizon.

To experimentally demonstrate this approach, we conducted a series of flight tests in a motion capture arena using a small, 790 g quadrotor MAV equipped with an ODROID-XU4 running the Robust EPC algorithm. The vehicle follows a high-speed straight line trajectory with artifical noise injected during the fastest portion, as shown in Figure 6. Figure 6 illustrates the resulting constraint satisfaction performance and it is clear that Robust EPC successfully satisfies constraints whereas EPC without the robust extension does not. For more details, see [13].

V. Conclusions and Future Work

Feature detection and tracking experiments with FLIR images and video sequences continue to show promise for performing visual-inertial state estimation with small, uncooled thermal cameras. Work in FY17 will aim to bring this to fruition with a complete, onboard implementation. Once that is complete, we expect that we will be able to demonstrate autonomous navigation and landing at night with a thermal camera. Work with our new predictive control frameworks has demonstrated their efficiency and accuracy in extensive experiments. In FY17 will will integrate this with both a short-range actual depth sensor and a long-range, passive structure from motion algorithm to realize our original scenario of autonomous approach to, detection of, and safe landing on structures representative of rooftop landings for persistent surveillance. Transition of these capabilities is underway to user agencies.

References

[1] V. R. Desaraju, N. Michael, M. Humenberger, R. Brockers, S. Weiss, and L. Matthies, "Vision-based Landing Site Evaluation and Trajectory Generation Toward Rooftop Landing," in *Proc. of Robot.: Sci. and Syst.*, Berkeley, CA, Jul. 2014.

[2] V. R. Desaraju, N. Michael, M. Humenberger, R. Brockers, S. Weiss, J. Nash, and L. Matthies, "Vision-based landing site evaluation and informed optimal trajectory generation toward autonomous rooftop landing," *Auton. Robots*, 2015.

[3] R. Brockers, M. Humenberger, S. Weiss, and L. Matthies, "Towards autonomous navigation of miniature UAV," in *IEEE Embedded Vision Workshop*, Columbus, OH, Jun. 2014.

[4] H. Hirschmuller, "Accurate and efficient stereo processing by semi-global matching and mutual information," in *IEEE Conf. on Computer Vision and Pattern Recognition.*, Jun. 2005.

[5] C. X. Guo, D. G. Kottas, R. C. DuToit, A. Ahmed, R. Li, and S. I. Roumeliotis, "Efficient visual-inertial navigation using a rolling-shutter camera and inaccurate time stamps," in *Proc. of Robot.: Sci. and Syst.*, Berkeley, CA, Jul. 2014.

[6] M. Li and A. Mourikis, "Vision-aided inertial navigation with rolling-shutter cameras," *Intl. J. of Robot. Research*, 2014.

[7] J. Hedborg, E. Ringaby, P.-E. Forssen, and M. Felsberg, "Structure and motion estimation from rolling shutter video," in *ICCV Workshops*, 2011.

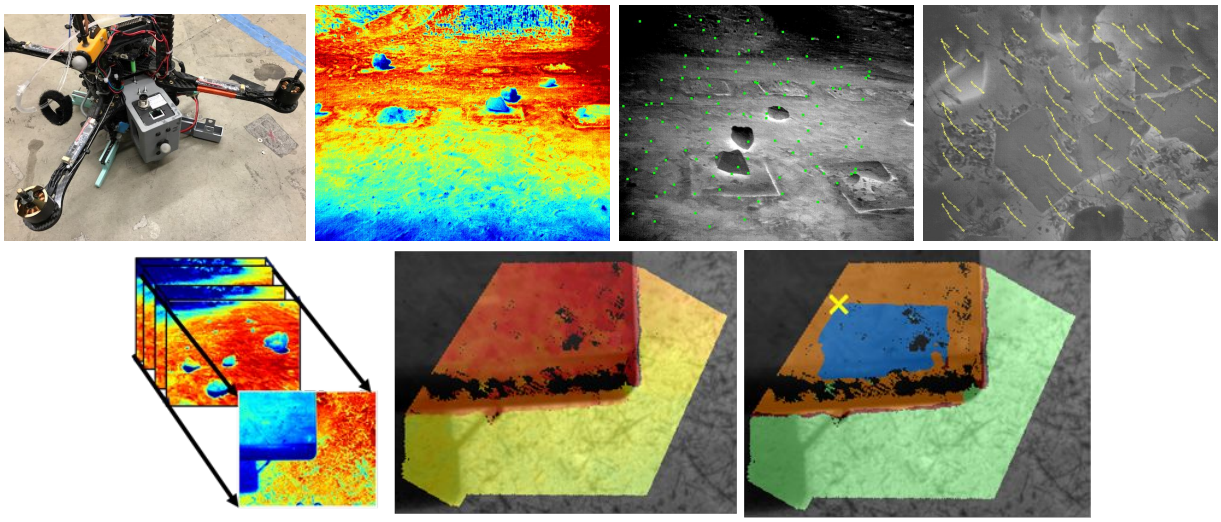


Fig. 4: (Top) Feature detection and tracking across a sequence of frames on a dataset collected using thermal-infrared camera. (Bottom) Qualitative results of landing site estimation using a thermal-infrared camera.

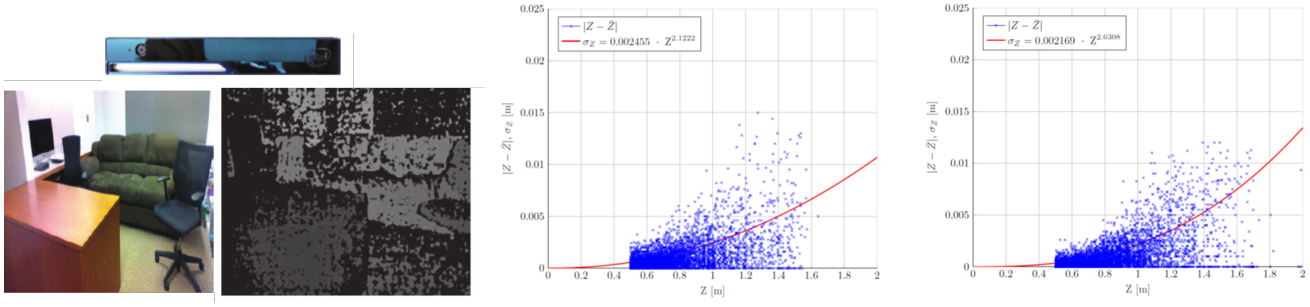


Fig. 5: (Left) RealSense and sample color and depth data. (Middle) Best curve fit of outdoor data during daytime and (Right) nighttime.

[8] D. Mitrovic, S. Klanke, and S. Vijayakumar, "Adaptive optimal feedback control with learned internal dynamics models," *From Motor Learn. to Inter. Learn. in Rob.*, vol. 264, pp. 65–84, 2010.

[9] A. Gijsberts and G. Metta, "Real-time model learning using incremental sparse spectrum gaussian process regression," *Neural Networks*, vol. 41, pp. 59–69, 2013.

[10] J. Engel, T. Schops, and D. Cremers, "LSD-SLAM: Large-scale direct monocular SLAM," in *European Conference on Computer Vision*, 2014.

[11] H. Bay, T. Tuytelaars, and L. Van Gool, "Surf: Speeded up robust features," in *European conference on computer vision*, Springer, 2006, pp. 404–417.

[12] J.-Y. Bouguet, "Pyramidal implementation of the affine lucas kanade feature tracker description of the algorithm," *Intel Corporation*, vol. 5, no. 1-10, p. 4, 2001.

[13] V. R. Desaraju, A. E. Spitzer, and N. Michael, "Experience-driven Predictive Control with Robust Constraint Satisfaction under Time-Varying State Uncertainty," in *Proc. of Robot.: Sci. and Syst.*, Boston, MA, Jul. 2017.

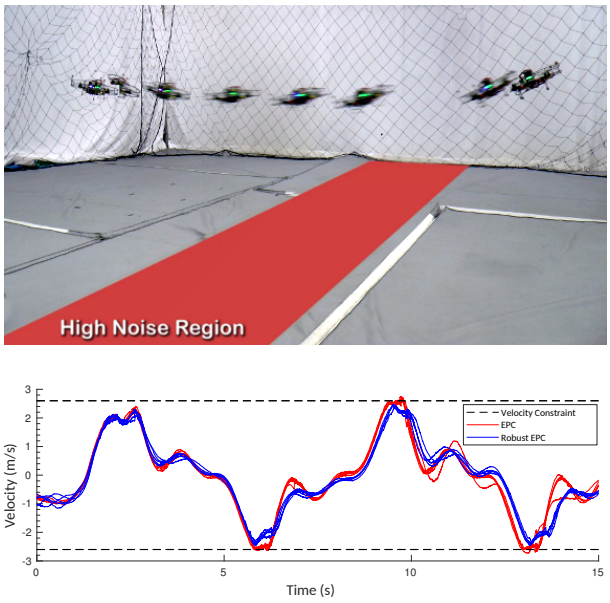


Fig. 6: (Top) The back and forth trajectory visualized using video stills. Region with injected noise is denoted in red. (Bottom) REPC satisfies velocity constraints in the presence of state-varying uncertainty.

Cooperative Aerial Reconnaissance with Airborne Sense and Avoid for Micro Air Vehicles

JX-16-17-4

Cevahir Cigla, Roland Brockers, and Larry Matthies
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA
{Cevahir.Cigla, brockers, lhm}@nasa.jpl.gov

Sarah Tang, Kartik Mohta, and Vijay Kumar
University of Pennsylvania
Philadelphia, PA
{sytang, kmohta, kumar}@seas.upenn.edu

***Abstract*—In order to enable teams of aerial robots to work close to each other in either cooperative or independent tasks, the robots need to be able to avoid hitting each other. All current approaches either use a centralized planner which knows the positions of the individual robots or use an external motion capture system to feed in the positions of neighbors to each of the robots. We have developed onboard perception and planning algorithms using stereo vision sensors in order to provide a self-contained solution for the short-range (5 to 20 meter) inter-robot collision avoidance problem.**

***Index Terms*—Cooperative Reconnaissance, Motion Planning, Stereo Vision based Perception, Collision Avoidance**

I. Introduction

A. Background and Technical Challenges

There has been research on several applications of teams of Micro Aerial Vehicles (MAVs) for cooperative missions, such as team reconnaissance and cooperative transportation of heavy payloads [1][2][3][4]. These applications require localization of all the aircraft in the team in order to plan the motions of each team member. In work to date, this localization knowledge has been provided from an external tracking system (e.g. Vicon, GPS) [3], with the aid of visual targets placed on each aircraft to simplify the vision problem for relative state estimation [4], or through multi-agent SLAM algorithms that use information about landmarks observed by multiple aircraft [5]. Multi-agent SLAM algorithms inherently require communication between the agents, or between the agents and a central mapping/planning system. In some circumstances, this communication requirement can be a drawback; at a minimum, it incurs the latency of the communication and map update cycles before knowledge of the motion of other aircraft is available.

There has been little to date work on using onboard perception systems for MAVs to directly detect other nearby aircraft. For large UAVs, the problem of detecting other aircraft to avoid collisions is called the Sense and Avoid (SAA) problem; there has been considerable work on this [6], but it does not scale down to MAVs, because the sensors are too big, the algorithms are too complex, or the systems are not sufficiently generalizable. There has been work on vision systems for MAVs to detect and track moving objects on the ground using a single onboard camera [7]; however, it is challenging to adapt this to perceiving other aircraft, because the absolute distance and velocity of other aircraft is unobservable with monocular imagery unless additional constraints can be applied, like known size of the other aircraft. There has been a substantial amount of work on autonomous navigation of ground robots among other moving objects on the ground, using range sensors like lidar or stereo vision for perception [9][10]. In principle, some of this work can be adapted to MAV SAA; however, the problems are not identical, since MAVs can move in full 3-D whereas ground vehicles stay on the ground. This could lead to different choices of representation and temporal update algorithms.

Multi-robot planning has become an active area of research. Broadly, there are two approaches for solving the multi-robot planning problem. The first is a coupled approach where the problem is solved by including all the robots in a single joint planner. These approaches typically guarantee optimality and completeness. However, as the search space grows exponentially with the number of robots, they quickly become computationally impractical. Decoupled approaches, as the name suggests, on the other hand plan for each robot independently. In this task, we are interested in decentralized planners that do not require explicit communication between the participating robots. The only way to achieve this is by using a decoupled approach.

In the aerial robotics field, obstacle avoidance has been addressed in many ways for static scenes [11][12][13] and in cooperating robot scenarios [14]. There has been some work [15][16] with non-cooperative robots, but it has only been demonstrated in simulation or using motion capture systems, and has not been done yet with real perception systems.

It is important to enable teams of MAVs to conduct cooperative or independent missions without requiring external tracking systems or artificial tracking targets on each MAV for detecting each other. The ultimate solution for this is likely to include a combination of multi-agent SLAM and onboard

direct perception of other nearby aircraft. The former problem is being addressed both within and outside the MAST program. The latter is not being addressed at all; it is important and timely to begin to tackle it.

B. Research Objectives

The Jet Propulsion Laboratory (JPL) and the University of Pennsylvania (UPenn) collaborated in a two-year effort to address airborne SAA for MAVs, using fully self-contained perception systems both for localization of each aircraft (e.g. [17]) and to perceive other nearby aircraft. We began by addressing basic maneuvers of two aircraft navigating close to each other; future work will generalize this to a larger team of robots performing a cooperative task where the inter-robot perception is used in distributing the task among the robots.

C. Technical Approach Summary

We assume that there is a global plan given to each of the robots in a team, which they are already following, but there is a need to avoid collisions with each other. Our approach to the SAA problem uses platforms with stereo cameras providing 3D perception coupled with temporal tracking to detect the in-air moving obstacles. The information from the perception system is fed to an estimator that continuously estimates the relative position and velocity of other nearby aircraft. The estimated state of all the tracked aircraft is then sent to a planner that perturbs the robot's originally planned path in order to avoid collisions with other agents in its vicinity. In this way, a robot is able to reactively avoid the collisions with dynamic obstacles without requiring explicit communication between the robots.

II. Related Work

Prior work outside the MAST program was discussed above. Within MAST, this research leverages and extends results from MAST tasks CNC-15.4.3 and SPP-15.1.3. Task CNC-15.4.3 developed a user focused multi-agent aerial system with advanced trajectory planning capabilities for autonomous and semi-autonomous surveillance. Task SPP-15.1.3 integrated dense stereo vision and optical flow to detect and avoid static obstacles. Here, we have extended this to detect and estimate the motion of other nearby moving MAVs. This task will extend and integrate results from the two previous tasks.

III. Approach

In this project, we leveraged and extended prior work done in the MAST program by JPL and UPenn. Our collaboration focused on two tracks. The first track, led by JPL, involved detection of moving obstacles around the robot using only local sensing. The second track, led by UPenn, extended the autonomy and robustness of the aerial swarm framework, QuadCloud [3], by integrating the obstacle detection capabilities developed by JPL into the planning sub-system with the goal of maintaining optimal area coverage.

The QuadCloud framework provides the base capabilities to control and navigate a team of aerial robots in a coordinated fashion. Currently, each robot in the QuadCloud system carries sufficient autonomy to do state estimation and control on board the robot. The robots have a downward facing camera in order to monitor the area beneath them. The operator can specify surveillance sites to the system and the system assigns goal positions to individual robots in the team and plans collision-free trajectories for each robot.

While moving along these trajectories, if there are external disturbances or estimation errors, the robots can be pushed too close to each other. To address this potential issue of collision among cooperative aerial agents operating in close proximity to each other, we require onboard detection and estimation of the motion of other nearby MAVs. At close range, this is most easily done using stereo vision for 3-D perception and object segmentation, because stereo vision gives the absolute position of other objects, together with some form of visual tracking to estimate motion of other objects. Onboard, dense stereo vision is available from our prior obstacle avoidance work under the MAST program [8]. Visual tracking of other moving objects can be achieved in several ways, including by tracking segmented objects in 2-D maps [10], or by tracking sparse features in images [20]. Since we segment 3-D objects from the stereo data in each frame, we perform tracking by data association of 3-D objects from frame to frame.

Our scenario is that a high-level planner provides trajectories to each robot, but the robots may come close to each other due to localization errors, external disturbances, or other factors. The trajectories of the robots need to be modified in real-time in order to avoid collisions with the other robot, while minimizing deviation from the pre-defined trajectory given by the global planner. Among the many local reactive planning schemes, the reciprocal collision avoidance approach [15] is the most suited to this problem since it requires no cooperation or explicit communication between the robots. During the first year, our focus was on a simple reactive collision avoidance problem with two aerial robots where the collision avoidance is accomplished by adding velocity perturbations to the trajectory from the high-level planner. The drawback with this approach is that it assumes that the robots have first-order dynamics and can instantaneously achieve a desired velocity. This assumption does not hold for most robots and leads to large control inputs to quickly change the robot's velocity. During the second year, we looked at reciprocal collision avoidance approaches that consider the higher order

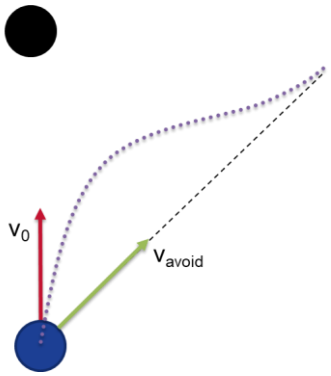


Fig. 1. The reactive collision avoidance algorithm computes a new velocity such that while changing the velocity direction the trajectory of the robot following higher-order dynamics avoids collision with obstacles.

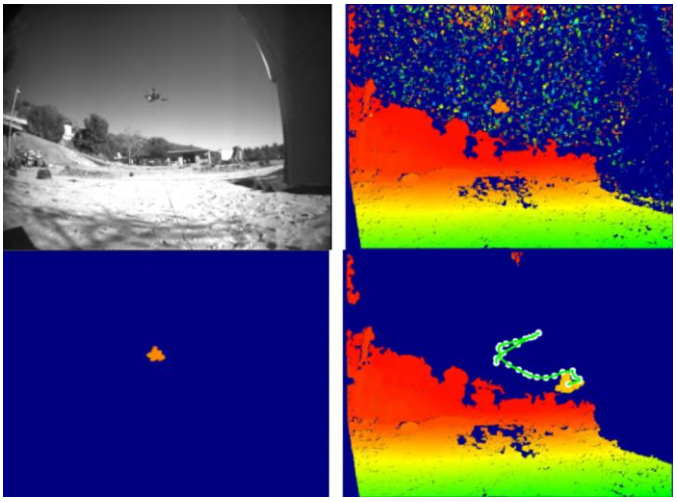


Fig. 2. Stereo-vision based blob tracking algorithm. a) Left stereo image, b) unfiltered stereo disparity map, c) filtered stereo disparity map, d) tracking the target quadrotor over time.

dynamics of the system and perform smooth collision avoidance maneuvers. We implemented the approach described in [21], which extends the original velocity obstacle concept to higher order systems. The primary idea in this approach is that it constructs a continuous control obstacle which represents the set of continuous control sequences which would lead to a collision between the robots. As shown in Figure 1, the algorithm computes a desired velocity such that the trajectory of the robot while changing the velocity direction avoids collision with the obstacles around it. This change in velocity is achieved using smooth changes in control inputs avoiding the problem of large control inputs caused by the standard velocity obstacle method. We continue to examine more deliberative planning approaches (for example, [22]), which, if able to run at high rates, may provide superior performance with noisy sensors.

The robots that were used for the QuadCloud project required changes for this project, since we need additional sensing for the sense-and-avoid system. Our current platform is the AscTec Pelican quadrotor with a stereo camera pair and Intel NUC computer mounted on it in order to allow us to run the sense-and-avoid system locally on the robot.

In our initial testing of the obstacle avoidance algorithms, we used a motion capture system for aircraft state estimation in order to focus on the planning and control issues. We shifted from the motion capture system to a visual-inertial estimation framework (e.g. [17]) as we integrated the visual tracking algorithms. All of the vision-based state estimation solutions, as they currently exist, lead to each robot having a separate reference frame, where typically the starting position of the robot is taken as the origin of the frame. Coordinated behavior would be improved by having a common reference frame for all the robots. GPS will provide a coarse estimate of this inter-robot frame transform, but the approach we develop for detecting moving objects can be used to refine this frame transform allowing us to plan and coordinate in a safer manner with lower tolerances. To simplify integration and testing, we also used a two-robot leader-follower scenario that did not require a common reference frame for both aircraft.

A. Research Accomplishments

During the first year of the project, we implemented a stereo-vision-based perception system to detect targets that are flying nearby. We first developed algorithms for detecting and tracking nearby quadrotors (“others”) from stationary stereo cameras, and then extended this to using stereo cameras onboard a moving quadrotor (“self”). Our approach computes 3D connected components from stereo depth maps, which correspond to small blobs within the disparity map, filters detected blobs based on their size and their position in the image, and tracks them over time (Figure 2). If a blob is of the size of the target quadrotor, is not located at the border of the image (which could otherwise e.g. be a large object mostly out of frame), and is tracked successfully over a number of successive frames, it is designated as a target. The blob tracker additionally uses a Kalman filter to estimate the target position

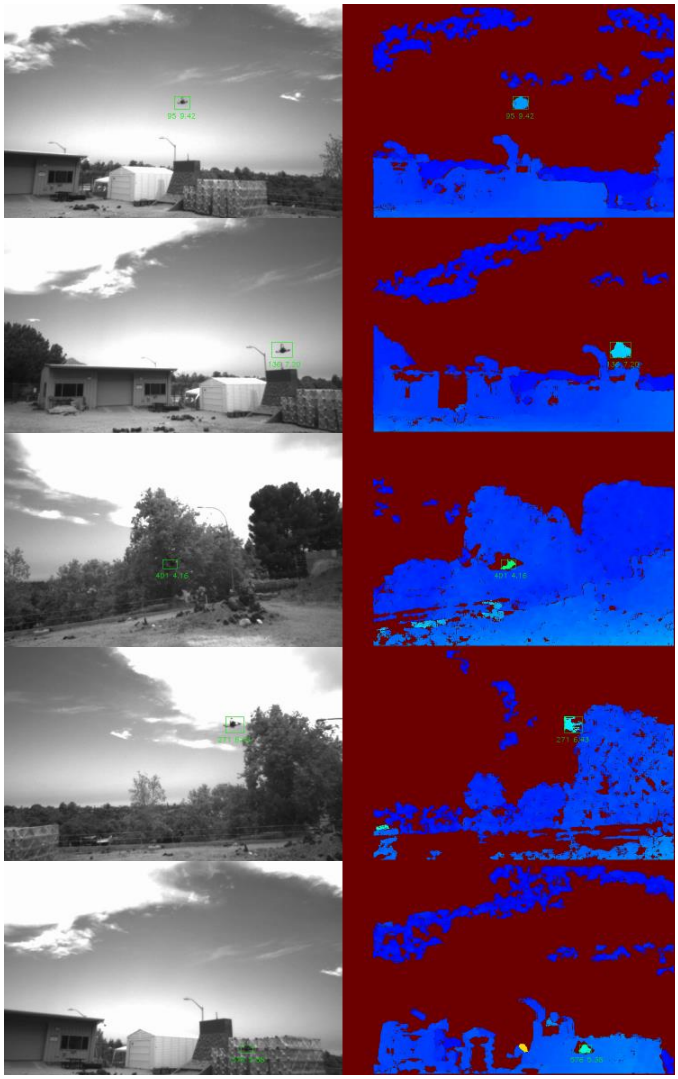


Fig. 3. Tracking of a flying quadrotor from moving “self” quadrotor; note that changes in the background scene reflect “self” motion. Distance to the “other” is around 10 meters. Tracking works both with the target silhouetted against the sky and viewed against clutter, like the tree in the third frame.

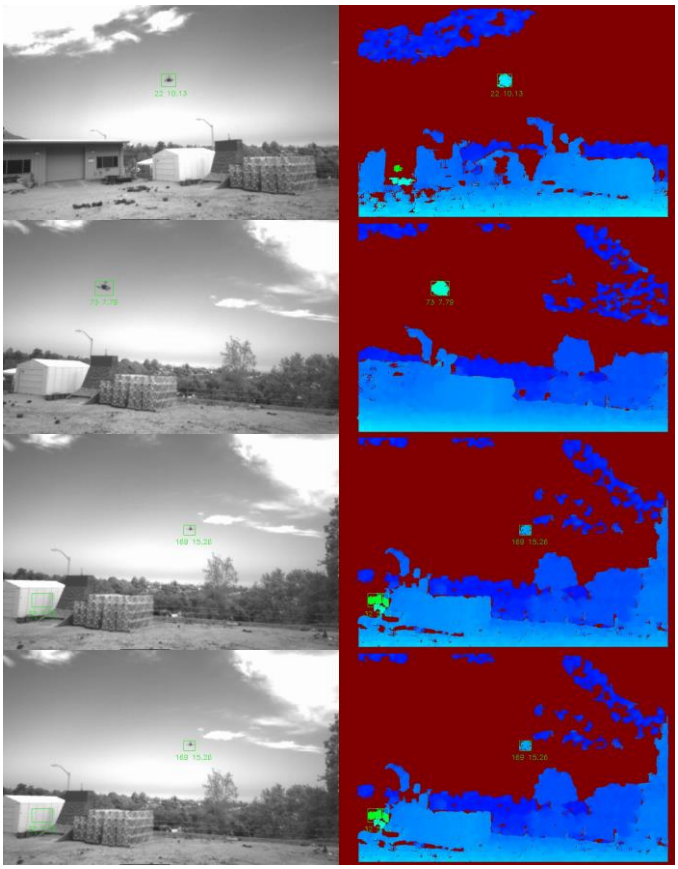


Fig. 4. Another tracking example, where distance to the other quadrotor ranges up to 15 meters (top to bottom).

and velocity in 3D, which allows predicting future target positions.

The states of "other" quadrotors (there can be more than one) are estimated in a fixed world reference frame. Onboard the "self" quadrotor, "self" state estimates are used to aid tracking of moving others from the moving camera platform, as well as to transform new position and velocity estimates of others into the world reference frame to update state estimates of the others. Some refinements to the detection and tracking algorithms have been added in the process, such adding the aspect ratio of detected blobs (in addition to diameter) in depth maps as a feature to help eliminate blobs that are not other quadrotors and to improve reliability of associating correct blobs from frame to frame for tracking. Small blobs very near each other in each depth map are merged, to improve detection in each frame, since individual other quadrotors sometimes appear as more than one disconnected blob in the depth maps. A "coast" mode has been added to the tracker to assist with reacquiring previously seen others if detection is missed for a small number of frames, for example due to errors in the depth maps, blob detection, or self-pose estimation. A typical tracking result is shown in Figure 3, where a flying quadrotor is tracked for 576 frames and then goes out of the field of view. The color image and the estimated disparity maps are illustrated with the tracked object in a green rectangle. The numbers correspond to the number of track frames and the distance of the object with respect to the camera. In the figure,

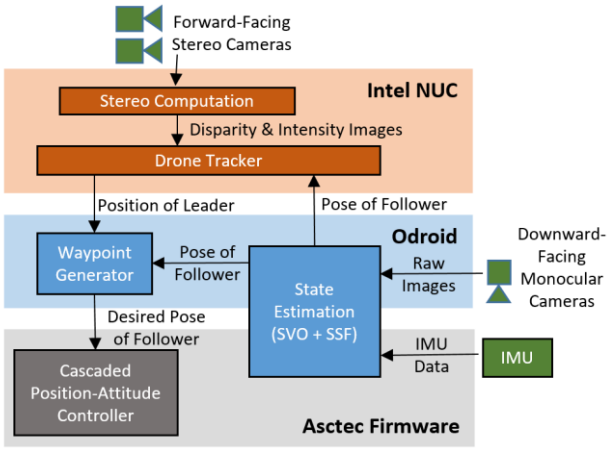


Fig. 5. System architecture for leader-follower scenario.

the quadrotor is tracked up to 10 meters away in front of a textured and untextured backgrounds. In Figure 4, another quadrotor is tracked up to 15 meters away, which also comes closer than 2.5 meters under fast camera motion. Thus, the stereo based tracking yields reliable results as long as accurate pose and disparity maps are provided.

This system currently runs on an AscTec Pelican quadrotor, with an Intel Core2Duo processor that generates depth maps from WVGA resolution stereo images (750x480) in about 0.25 seconds/frame, using a SAD1 local block matching (LBM) stereo algorithm, and performs tracking in about 0.02 seconds.

Although depth maps are usually fairly good, it was initially observed that occasional frames of bad data occurred, which was due to a vibration issue affecting transient vertical alignments of the stereo camera pair during outdoor flights where camera exposure times are very small. After designing a new camera mount that dampens the vibrations that go into the stereo camera mount, we did not observe any further issues with disparity map dropouts.

We also conducted off-line experiments with the semi-global block matching (SGBM) stereo algorithm in this application. SGBM has potential to generate denser depth maps, which may aid detection and tracking, but at higher computational cost. SGBM and other more powerful stereo algorithms are options for future extensions, as hardware acceleration for this class of algorithms becomes available.

To demonstrate the capability of inter-robot collision avoidance we implemented a leader-follower behavior using two quadrotors. The leader quadrotor was manually piloted and the follower robot autonomously follows the leader while maintaining a user-define safety distance from it.

Figure 5 shows the pipeline of our implementation of the leader-follower scenario. The stereo images are received on the Intel NUC and a disparity image is computed. Our tracker computes the position of the leader using the intensity and disparity images along with the pose estimates of the follower that are obtained from the state estimation module based on Single Sensor Fusion (SSF) [23] and Semi-Direct Visual Odometry (SVO) [19] algorithm that uses the images from the downward facing camera.

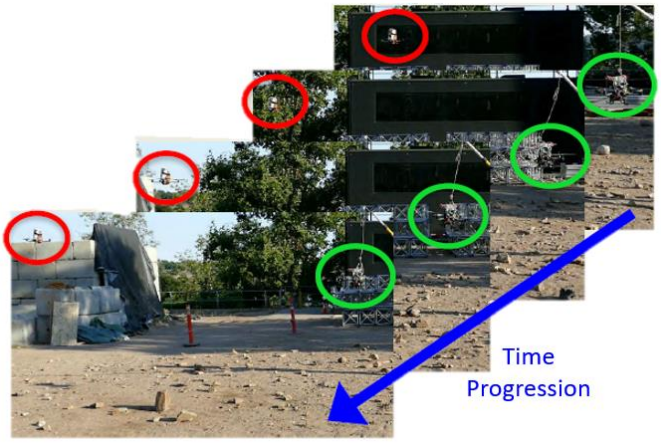


Fig. 6. Follower drone (green) autonomously tracks a manually controlled leader drone (red)

The waypoint generator uses the position estimates of the leader and the follower to compute the desired position of the follower, that is then sent to the AscTech firmware. This desired position is calculated by moving 25% along the line joining the center of the follower to the center of the leader. Further, if the calculated desired position lies within the safety distance threshold of the leader, then the desired position is set to the point of intersection of the line with a sphere center about the leader with a radius equal to the safety distance threshold. Finally, the AscTech firmware runs a cascaded position and attitude control loop that takes the follower drone to the desired position. Outdoor test results of a leader-follower experiment are shown in Figure 6.

In parallel to the perception module being developed at JPL, UPenn developed and tested different approaches for reactive collision avoidance. The high-level plan to the system is provided by a global planner, like the one used in the original QuadCloud system, or can be given by hand if desired. As the robots are following the trajectory from the high-level planner, the reactive collision avoidance algorithm takes care of shifting the trajectories of each robot away from obstacles in the vicinity. Some initial simulation results are shown Figure 7. The initial trajectory generated by the planner for the green robot is a dynamically feasible and smooth polynomial trajectory. As the robot starts moving along the trajectory and detects that there is another robot near it, the reactive planner adds a perturbation to the trajectory to go around the other robot. Here no explicit communication is required between the robots and the robots are not explicit cooperating to achieve collision avoidance. In order to test the collision avoidance algorithm on real robots, we used a motion capture system to simplify the robot estimation and detection tasks. Figure 8 shows snapshots from an experiment where the two robots are commanded to exchange positions while doing the collision avoidance. Figure 9 shows the plots of the position of the two robots as they avoid collision with each other while following different trajectories.

UPenn has also developed a centralized multi-robot planning and trajectory generation method that explicitly accounts

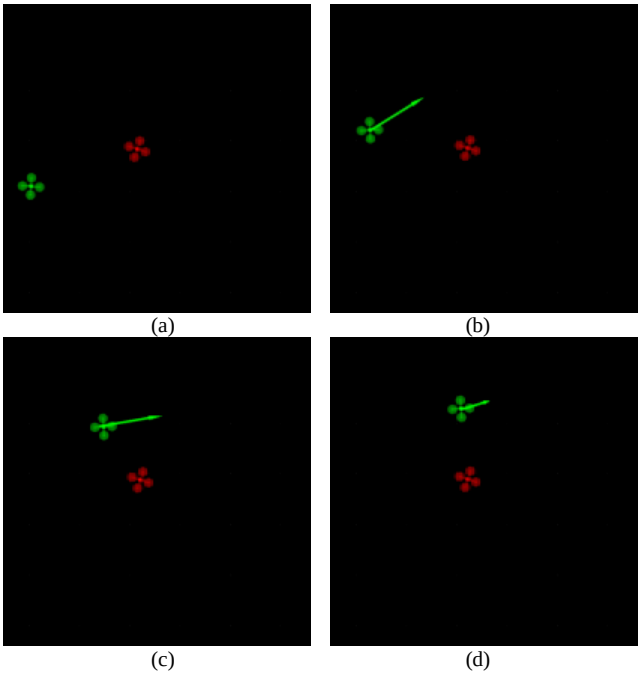


Fig. 7. Simulation of a reactive collision avoidance algorithm where the red quadrotor is just hovering at the middle while the green quadrotor is sent a goal towards the right and has to maneuver around the red one.

for higher-order dynamics, but requires communication and exchange of state estimates between the robots [22]. In experiments, we have demonstrated successful execution of solution trajectories and verified the robustness of the robots' trajectory tracking to unmodeled external disturbances, including the aerodynamic interactions between co-planar neighbors. We have also extended the algorithm to three dimensions where quadrotors navigate away from the downwash of their neighbors to improve safety in three dimensional workspaces.

B. Metrics for Performance Evaluation:

To date, we have been conducted separate evaluation the perception and collision avoidance algorithms.

The perception system works well if the target is seen in front of a low-texture background (e.g. sky) and also tracks when the target is in front of clutter below the horizon; tracking may be lost briefly as the target crosses the horizon, but is reacquired a few frames later. Target position and velocity tracks are generated at about 4 Hz at WVGA image resolution on our Asctec Pelican test platform, which is limited by the frame rate of our current stereo vision system. Experiments showed that at this resolution tracking is robust up to 15 m distance.

For the collision avoidance algorithm, we have tested it in simulation and obtained reliable performance up to speeds of 3 m/s between the robots. We also demonstrated successful live operation with two robots in the motion capture environment. Integration with the perception system was deferred due to flex in the stereo camera mount on some of the robots, which degrades stereo vision performance. This is an engineering issue that can be fixed by fabricating new camera mounts.

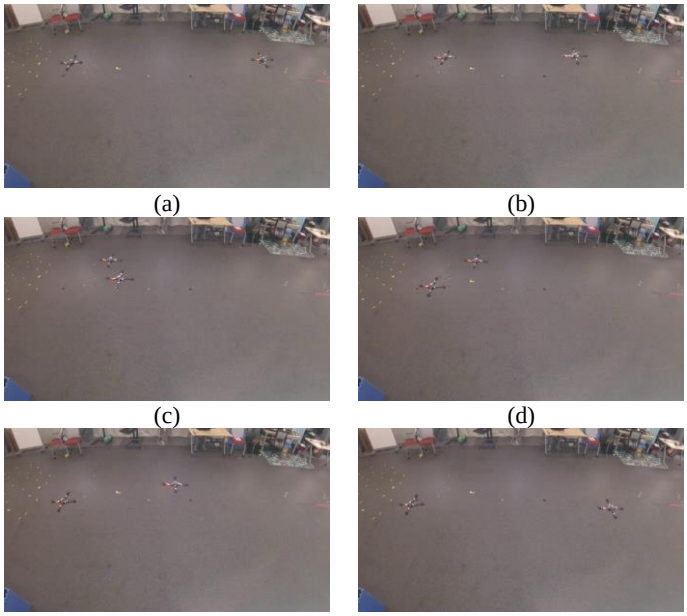
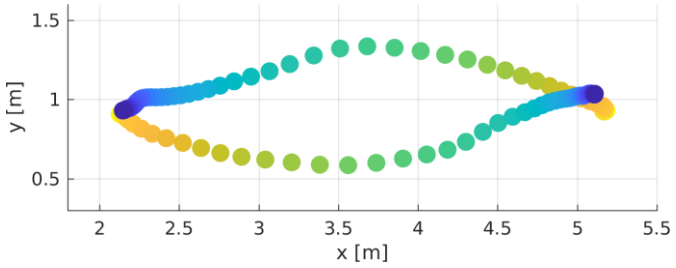


Fig. 8. Snapshots from an experiment where the robots are commanded to exchange positions while running the collision avoidance algorithm. In this experiment, the motion capture system was used for providing the positions and velocities of the robots.

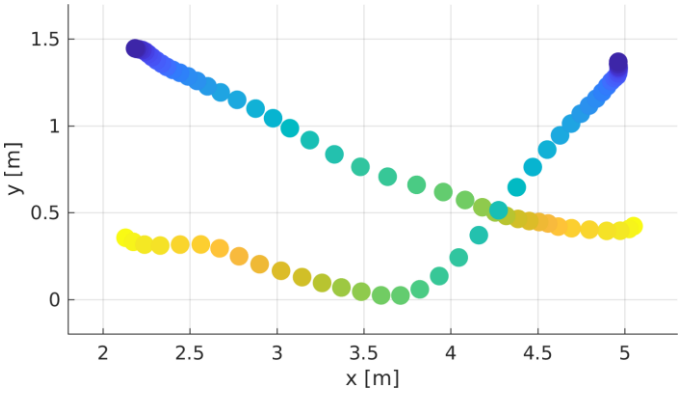
IV. Conclusions and Further Work

Experiments with the onboard implementation of the perception algorithms have demonstrated that using stereo vision for detection and tracking works well at ranges of up to about 15 m with the resolution and frame rate of our current vision system. This distance can increase with higher resolution stereo vision systems using ASICs for computation. The approach to planning works well in simulation and in live testing in a motion capture arena. Fully integrated experiments were conducted successfully for a leader-follower scenario. Fully integrated experiments with cooperative collision avoidance were postponed due to inadequate rigidity in camera mounts on some of the required testbed aircraft; this is readily solvable in future work.

The stereo vision-based detection and tracking algorithm developed here has potential use in the opposite problem of intercepting intruder aircraft for counter-UAS (unmanned air systems) applications.



(a) Robots exchanging positions



(b) Robots moving along an X pattern

Fig. 9. Plots showing the position of the robots from experiments when the robots are commanded to (a) exchange their positions and (b) go to waypoints leading to them crossing in an X pattern. The time progression is denoted by the changing colors of the disks going from yellow at start to blue at the end.

References

[1] A. Renzaglia, L. Doitsidis, A. Martinelli and E.B. Kosmatopoulos, "Multi-Robot 3D Coverage of Unknown Areas," The International Journal of Robotics Research, 2012

[2] M. Schwager, D. Rus and J. J. Slotine, "Decentralized, Adaptive Coverage Control for Networked Robots," International Journal of Robotics Research, March, 2009, vol. 28, no. 3, pp. 357-375.

[3] K. Mohta, M. Turpin, A. Kushleyev, D. Mellinger, Nathan Michael, and V. Kumar, "QuadCloud: A Rapid Response Force with Quadrotor Teams", International Symposium on Experimental Robotics, 2014

[4] K. Sreenath and V. Kumar, "Dynamics, Control and Planning for Cooperative Manipulation of Payloads Suspended by Cables from Multiple Quadrotor Robots", Robotics: Science and Systems, 2013

[5] J. Dong, E. Nelson, V. Indelman, N. Michael, F. Dellaert, "Distributed real-time cooperative localization and mapping using an uncertainty-aware expectation maximization approach," IEEE International Conference on Robotics and Automation (ICRA), 2015

[6] D. Dey, C. Geyer, S. Singh, M. Digioia, "Passive, long-range detection of aircraft: towards a field deployable Sense and Avoid system," Field and Service Robotics Conference, 2009

[7] D. Meier, R. Brockers, L. Matthies, R. Siegwart, S. Weiss, "Detection and Characterization of Moving Objects with Aerial Vehicles using Inertial-Optical Flow", International Conference in Intelligent Robots and Systems (IROS), 2015

[8] L. Matthies, R. Brockers, Y. Kuwata, S. Weiss, "Stereo vision-based obstacle avoidance for micro air vehicles using disparity space," IEEE International Conference on Robotics and Automation (ICRA), 2014

[9] A. Talukder, L. Matthies, "Real-time detection of moving objects from moving vehicles using dense stereo and optical flow," IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2004

[10] M. Bajracharya, B. Moghaddam, A. Howard, S. Brennan, L. Matthies, "A fast stereo-based system for detecting and tracking pedestrians from a moving vehicle," International Journal of Robotics Research, 2009

[11] L. Matthies, R. Brockers, Y. Kuwata, S. Weiss, "Stereo vision-based obstacle avoidance for micro air vehicles using disparity space," IEEE International Conference on Robotics and Automation (ICRA), 2014

[12] F. Kendoul, “Survey of advances in guidance, navigation, and control of unmanned rotorcraft systems,” Journal of Field Robotics, 29(2), 2012

[13] S. Liu, M. Watterson, S. Tang and V. Kumar, "High speed navigation for quadrotors with limited onboard sensing," 2016 IEEE International Conference on Robotics and Automation (ICRA), Stockholm, 2016, pp. 1484-1491.

[14] M. Turpin, N. Michael, and V. Kumar, “CAPT: Concurrent assignment and planning of trajectories for multiple robots,” The International Journal of Robotics Research, vol. 33, no. 1, pp. 98–112, 2014.

[15] Bareiss, D., van den Berg, J., "Reciprocal collision avoidance for robots with linear dynamics using LQR-Obstacles," in IEEE International Conference on Robotics and Automation (ICRA), 2013

[16] J. Alonso-Mora, T. Naegeli, R. Siegwart, “Collision avoidance for aerial vehicles in multi-agent scenarios,” Auton Robot (2015) 39: 101-121.

[17] S. Weiss, M. W. Achtelik, S. Lynen, M. C. Achtelik, L. Kneip, M. Chli, R. Siegwart, “Monocular Vision for Long-term Micro Aerial Vehicle State Estimation: A Compendium.” J. Field Robotics, 30:803–831. doi:10.1002/rob.21466, 2013

[18] Klein, G., Murray, D., "Parallel Tracking and Mapping for Small AR Workspaces," in 6th IEEE and ACM International Symposium on Mixed and Augmented Reality, 2007. ISMAR 2007, pp.225-234, 13-16 Nov. 2007

[19] Forster, C., Pizzoli, M., Scaramuzza, D., "SVO: Fast Semi-Direct Monocular Visual Odometry," in 2014 IEEE International Conference on Robotics and Automation (ICRA), vol., no., pp.15-22, May 31 2014-June 7 2014

[20] D. Meier, R. Brockers, L. Matthies, R. Siegwart, S. Weiss, “Detection and characterization of moving objects with aerial vehicles using inertial-optical flow,” IEEE/RSJ International Conference on Intelligent Robots and Systems, 2015.

[21] M. Rufli, J. Alonso-Mora and R. Siegwart, “Reciprocal Collision Avoidance With Motion Continuity Constraints,” in *IEEE Transactions on Robotics*, vol. 29, no. 4, pp. 899-912, Aug. 2013. doi: 10.1109/TRO.2013.2258733

[22] S. Tang and V. Kumar, "A Complete Algorithm for Generating Safe Trajectories for Multi-Robot Teams," in International Symposium on Robotics Research (ISRR), Sestri Levante, Italy, September 2015.

[23] Weiss, Stephan, et al. "Real-time onboard visual-inertial state estimation and self-calibration of mavs in unknown environments." Robotics and Automation (ICRA), 2012 IEEE International Conference on. IEEE, 2012.

V. 2016 – METRICS

MAST FY16 ANNUAL REPORT METRICS TABLE	
Awards (Miscellaneous Recognitions and Awards)	0
Refereed/Journal papers	0
Conference presentations/proceedings	0
Joint Papers with ARL	0
Invention disclosures	0
Patent applications filed	0
Patents awarded	0
PhDs Award	0
Masters Awarded	0
Post Docs Supported	1
PhD Degree Students Supported	2
Master Degree Students Supported	0
Undergraduate Students Supported	1
CTA Member rotations into ARL	0
Visiting Researchers from ARL to the Member Sites	0
Employees hired by ARL from CTA students supported	0

Integrating Autonomous Navigation on MAST-scale Rotorcraft

JX-16-17-5

Roland Brockers
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA
brockers@nasa.jpl.gov

Vijay Kumar and Kartik Mohta
University of Pennsylvania
Philadelphia, PA
kumar@seas.upenn.edu
kmohta@seas.upenn.edu

***Abstract*— Many novel autonomous navigation algorithms have been developed during the course of the MAST program, including algorithms for visual-inertial state estimation, trajectory generation, and high-level functions like collision avoidance, autonomous exploration, collaborative lifting of heavy payloads, and autonomous roof-top landing for surveillance applications. While most of the developed technology was demonstrated on rotorcraft aerial vehicles weighing under 1 kg, it is still important to eventual users to get these capabilities into much smaller aircraft. In this task, we take MAST technology for autonomous flight to a vehicle that is significantly smaller than previously demonstrated, approaching for the first time palm-sized aerial vehicles capable of fully autonomous flight.**

***Index Terms*— Miniature, Micro Air Vehicle, Autonomous Navigation, MAST-scale**

I. Introduction

A. Background and Technical Challenges

The Army has a strong interest in micro or nano Unmanned Aerial Vehicles (UAV) to provide immediate tactical reconnaissance capability to small dismounted units; for example, such pocket sized systems are part of the long-term vision in the Army's Robotics and Autonomous Systems Strategy. Such UAVs need autonomous navigation capability for non-line-of-sight (NLOS) operation, GPS-denied operation, operation in confined spaces, and operation where high-bandwidth teleoperation is impractical. The MAST program has developed relevant sensing, processing, and algorithm capabilities, but integration and testing of these has largely taken place on relatively large UAVs of 0.5kg and more, or has not taken place yet at all.

With the advent of more powerful embedded processing units, and the increased maturity of vision algorithms, GPS denied navigation became possible for lighter weight, more power constrained vehicles [7][8]. For example, past MAST efforts showed that visual-inertial navigation can be done on a 12g processor board (hosting an Exynos4 system on a chip (SoC)) [1]. By pairing such a processing platform with a small, ultra-light-weight cell-phone camera (< 2g) and a light-weight MEMS IMU (< 1g), a complete avionics package that weighs below 25g becomes possible (Figure 1). With such a platform, a miniature UAV (e.g. Figure 2) can be accurately navigated in GPS denied areas.

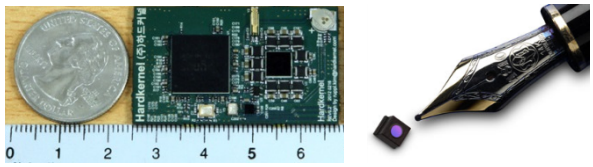


Fig. 1. Examples of small avionics. Left: Flight computer using a small 1.7 GHz quad-core ARM Cortex-A9 board, (12g). Right: CameraCube (< 2g).

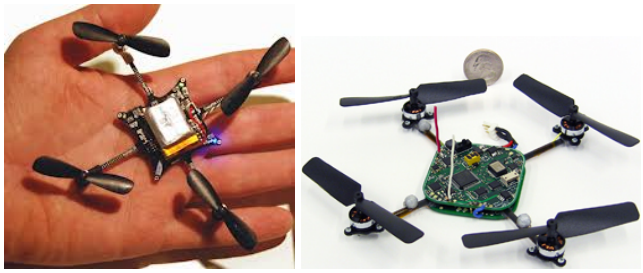


Fig. 2. Examples of small aircraft. Left: Crazyflie-nano (19g) [Crazy 2014]. Right: KMel nanoQuad (75g).

Similarly, with the necessary improvements, past MAST efforts on obstacle avoidance [3], safe landing [4][5], and trajectory generation [6], can be implemented on such a light-weight processor board and sensor suite. This yields a completely GPS independent, pocket sized UAV capable of autonomous, vision-based, indoor/outdoor state estimation, obstacle avoidance, and landing for power conservation and persistent surveillance. Leveraging other miniaturized sensors and/or processing elements from the Microelectronics Center, such as optical flow or stereo vision ASICs, might further contribute to these goals when these modules reach sufficient maturity.

B. Research Objectives

There are multiple candidate autonomous navigation technologies available for the above integration effort, and

multiple candidate UAV platforms to use for it. Leveraging a survey that was conducted in a supplemental effort in FY15, we selected a target hardware platform including airframe, sensors, and processing platform, and created an integration plan to enable GPS independent autonomy on a nano-scale UAV. More specifically, we seek to integrate the autonomy core-capabilities: accurate, high-rate, low level control; uninterrupted high-fidelity state estimation and self-calibration of the system; autonomous obstacle avoidance; safe landing; and optimal trajectory planning and following. The objective is to improve the algorithms enabling these capabilities on highly payload and computationally constrained vehicles through rigorous analysis on computational bottlenecks and corresponding algorithmic advancements. We will track progress on relevant modules from the Microelectronics Center and consider their integration if appropriate. The final outcome will be the smallest UAV in its category featuring GPS independent autonomy.

C. Technical Approach Summary

Based on a market survey conducted in a MAST supplemental task, a target miniature airframe was selected that is capable of carrying COTS sensing and processing hardware needed to execute current MAST autonomous navigation algorithms. The target size is a vehicle slightly larger than palm-size with a sensing and processing suite that weighs <50g. Core capabilities to be integrated on this miniature vehicle include GPS-denied, visual-inertial state estimation, on-board trajectory generation and following, visual safe landing, and, if the computing budget allows, visual obstacle avoidance.

II. Related Work

A. Synergy with Existing Work and External Collaborations:

This task combines prior work at JPL and UPenn on autonomous aerial vehicle navigation.

In the last decade, we have observed the introduction of a new generation of MAVs with footprints less than 20cm and masses below 200 grams. These include flapping wing vehicles like the Nano Hummingbird [12], DelFly Explorer [13] and the Robobee [14], fixed wing aircraft like the Black Widow [15] and rotorcraft like the Black Hornet [16]. Among commercially available MAVs, specifically quadrotors, the Crazyflie [2], Ladybird and ProtoX [17] are the most popular vehicles. However, these either suffer from poor battery life, can carry little to no payload, or are closed source, preventing access to low level APIs necessary for research.

JPL leverages previous work in map based visual-inertial state estimation for GPS independent UAV navigation [7][8] and map-free Inertial-Optical Flow (IOF) based state estimator capable of fast vehicle deployment (i.e. throw-and-go) [9], autonomous vision-based obstacle avoidance [3], and safe landing site detection [5][6]. UPenn leverages previous work in advanced trajectory planning for autonomous and semi-autonomous surveillance and highly accurate UAV low level control. This makes full use of the agility of the platform while still navigating on energy optimized trajectories.

III. Approach

We leverage a survey conducted in a MAST supplemental task in FY15 to quantify available UAV platforms and hardware and software solutions for sensing, processing, and communication from within MAST or other government programs, or as COTS products, with emphasis on open source hardware or software.

To include visual-inertial state estimation, visual obstacle avoidance, visual safe landing, and on-board trajectory generation as autonomous navigation capabilities on a MAST sized aerial vehicle, small cameras are key to the miniaturization objectives. Thus, we particularly focus on the use of smartphone-scale cameras and camera interfaces. Key COTS processor options include the Hardkernel Odroid boards already used in the MAST program, which carry a Samsung Exynos SoC, and the Intrinsyc Snapdragon Flight board that is specifically designed with UAS applications in mind and hosts a Qualcomm Snapdragon system on a chip (SoM). For the aircraft itself, we consider COTS multi-copters, in-house builds similar to previous KMel vehicles, and other custom build approaches recently identified by ARL and other MAST performers. Based on a study performed in the FY15 supplemental task, we laid out a conceptual design and plan at the beginning of this task for conducting this integration work in the remaining two years of the MAST program.

Although the focus of this task is not hardware development or aerial platform and sensor-suite development, we initially integrate the selected sensor suite on a candidate nano UAV. While we consider interfacing to sensors like an inertial measurement unit, magnetometer, air pressure sensor or even GPS to be trivial, interfacing the miniature camera will need more integration effort. Cell phone type cameras generally use a MIPI interface, which is often not directly available in COTS system on chip processing boards. Additionally, we anticipate small hardware adaptations in order to correctly map the system's control data output to the motor controller of the platform. Since the motor control loop on an inherently unstable multicopter platform requires high frame rate and real-time feedback, an additional low-level micro controller may be required specifically dedicated to this low-level control task.

The rest and main part of this task then focuses on the on-board implementation and improvement of the autonomy algorithms. The lead of this effort is split between JPL and UPenn based on their area of expertise. As a first step in the control and path planning area, UPenn implements low level attitude, velocity, and position control algorithm on the low level embedded processor leveraging prior work in this domain and adapt existing algorithms to the selected processor on the nano UAV. We seek to reduce the computational complexity of these low-level algorithms to enable high-rate, real-time closed loop control. At the same time, less computational complexity for the control tasks will allow implementing parts of the state estimation algorithms such as the state integration on the same embedded real-time processor, which allows a

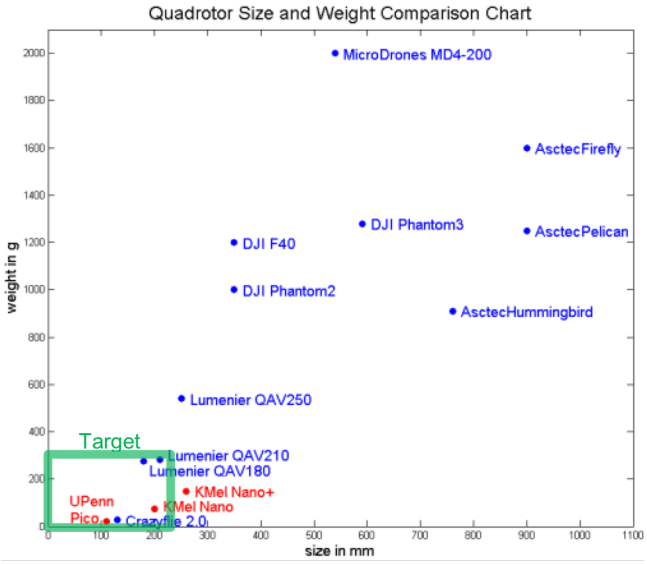


Fig. 3. Size weight comparison chart for quadrotors available at project start.

real-time and thus more accurate propagation step in the state estimation module.

The low level control is closely linked with the trajectory generation and path planning module. In a second step, UPenn implements a trajectory module on the main processing unit of the platform. Nano UAVs typically have a very low endurance. Thus, it is important that the trajectory and path planning module not only generates feasible paths from the system dynamics point of view but also paths that are minimal in energy for increased system endurance. We draw again from prior work to generate these trajectories to improve the algorithms in terms of runtime and complexity on the nano UAV. In particular, we seek to simplify the approach to both maximize performance and minimize computational complexity. The implemented solution then interacts with the vehicle on-board state estimation, obstacle avoidance, and landing module to design mission specific maneuvers and paths.

JPL leads the effort on the implementation and improvement of all vision-based modules including the visual-inertial state estimation, obstacle avoidance, and landing module. In a first step, an existing visual-inertial state estimation module is ported to the nano UAV. And as before, we seek to introduce selected algorithm simplifications such as improved key-frame selection reducing the number of stored key-frames while maintaining the same accuracy for the underlying map to reduce computational cost of the algorithms used. The estimator may be adapted to function with a front looking camera instead of a down looking one as it is currently implemented. This allows the simultaneous use of a camera for monocular state estimation as well as stereo-based obstacle avoidance in conjunction with a second front looking camera. We additionally plan to include a task scheduler to allocate specific system resources to each of the modules and sub-modules. This ensures sufficient processing power for accurate

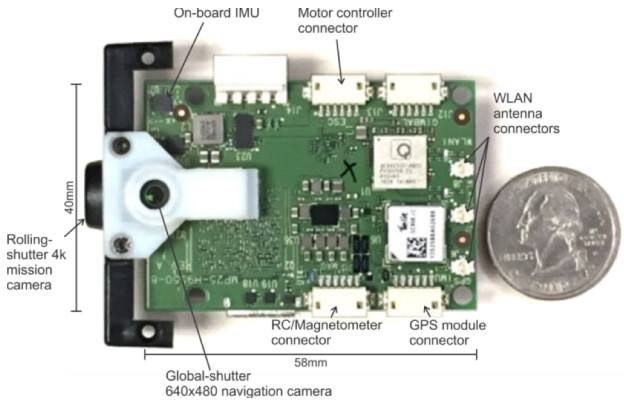


Fig. 4. Intrinsyc Snapdragon Flight board.

state estimation at all times. Depending on the available system resources, we will port in a similar way the fail-safe fall-back option using inertial cues and optical flow for state estimation to bridge failures of the map-based approach.

In a next step, autonomous landing is added to enable the vehicle to safely land for energy conservation, persistent surveillance, and potentially to re-charge for subsequent missions. We will implement a safe landing site detection algorithm that has been developed previously in the MAST program. Since frame rate 3D reconstruction for the selection of safe landing sites is not required for safe and stable control of the UAV, the landing site detection software can effectively run at a lower frame rate, minimizing computational requirements for this module.

Depending the computing throughput remaining after integrating the above modules, the an obstacle-avoidance module to the system is added to the system, for which a light-weight, forward-looking stereo rig provides visual perception. The development and implementation of an overarching feature extraction and matching module will eliminate redundant low level vision processing needed for all three vision modules (state estimation, obstacle avoidance, and landing). This low-level vision module will be designed in a modular fashion such that it can further be re-used in other MAST tasks that need visual feature information. For the obstacle detection and avoidance, JPL leverages and further improves its c-space representation in image space for efficient collision checking to minimize the computational load of this module while at the same time delivering high quality information to the interacting path planning module.

Optical flow and stereo vision ASICs under development in the Microelectronics Center have potential to contribute to this architecture, if they reach sufficient levels of maturity in time and can be integrated in a sufficiently cost-effective manner. If such hardware accelerators become available during the course of this task, we plan to adapt plans for this task as appropriate to include this technology on the miniature platform.

IV. Results

A. Research Accomplishments

At the beginning of this task, we surveyed commercially available solutions for sensors, processing platforms, and airframes for miniature rotorcraft UAV (Figure 3). Based on this survey, we selected the Intrinsyc Snapdragon Flight processing board, which offers sufficient processing power to execute our autonomous navigation algorithms. It also integrates all necessary sensors in one small module the size of a credit card. The board weights less than 20g and hosts a Qualcomm Snapdragon 801 SoC (Figure 4). The 801 features an ARM-based 2.26 GHz quad-core CPU, running a Linaro Linux kernel, and a DSP (QDSP6 V5A - 801 MHz+256KL2) for real time flight control. Integrated sensors include an on-board Invensense MPU-9250 IMU, a Bosch BMP280 barometric pressure sensor, a down-facing navigation camera (Onmivision OV7251, global shutter, 640x480), and a forward facing 4k camera module (Sony IMX135, rolling shutter).

We integrated the Snapdragon Flight board on a Lumenier QAV180 airframe, we selected as our main test vehicle (Figure 5) for implementing our autonomous navigation algorithms on the Snapdragon Flight board. The total weight of the vehicle is 320 g without battery, which gives us a flight time of about 5-10 minutes.

Figure 6 illustrates the system architecture on the Snapdragon Flight board. One particularity of the board is that it does not include PWM interfaces to communicate with the electronic speed controllers (ESCs) that drive the motors of the quadrotor. Instead, the communication interface is implemented digitally using I2C with a proprietary protocol that is not compatible with commercially available ESCs capable of communicating through I2C. As a result, we implemented a new I2C protocol within the PX4 flight stack on the DSP that allows outputting motor speed commands to the ESCs we selected for the QAV180 (Afro ESC), which use open source firmware that included a protocol for I2C communication.

After verifying the correct functioning of the setup in manual control, we fused position estimates from our VSLAM framework [1] (PTAM) with IMU readings using the PX4 sensor fusion filter.

Due to the split architecture of the Snapdragon Flight board (Figure 6 & 8), we need a way to communicate between the high-level code running on the ARM cores with the PX4 flight

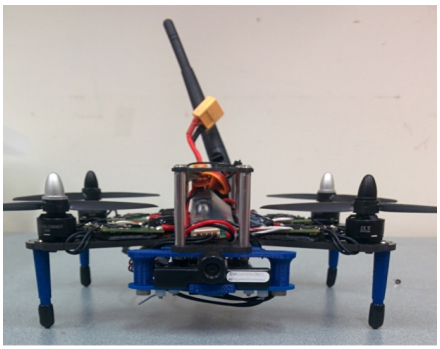


Fig. 5. Assembled Lumenier QAV180 with attached Snapdragon Flight board.

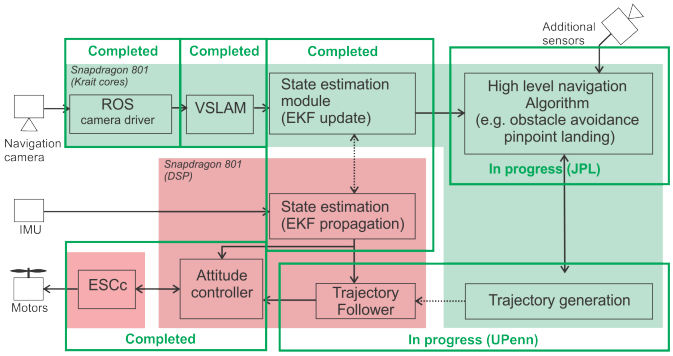


Fig. 6. System architecture and current software status on the Snapdragon Flight board.

stack running on the DSP, which uses the Mavlink protocol [11] for communicating with outside nodes running on the ARM cores. On the Snapdragon flight board, this interface is accessible through a local UDP port. Since all of our high level nodes use the ROS framework for communication between each other, we use the Mavros package [12] to send commands and receive data from the low-level controller and estimator running on the DSP.

The vehicle was tested in this setup in a laboratory environment with successful flights, where the vehicle autonomously hovered in place using pose estimates from the VSLAM algorithm.

In a next step, we replaced the PX4 state estimation filter with our own visual-inertial fusion filter. This filter combines visual odometry output from the vision system with IMU data in an Unscented Kalman Filter (UKF). A UKF is used since the quadrotor system with the IMU as the input for the prediction step makes the system nonlinear and compared to the more traditional EFK, the UKF has shown to have better accuracy with only a minor increase in computational cost. The data flow diagram of the UKF is shown in Figure 7.

The position, velocity and orientation estimates from the UKF are sent to the position controller which takes in the desired inputs from the trajectory generator and produces desired thrust and orientation commands. These commands are then sent to the attitude controller running on the DSP which in turn sends the motor commands to the speed controllers (Figure 8).

The commands for the position controller are provided by a trajectory generation module which takes in a series of waypoints and the limits on the dynamics of the robot and solves for a smooth polynomial trajectory using a Quadratic Program (QP). We can add linear constraints representing the

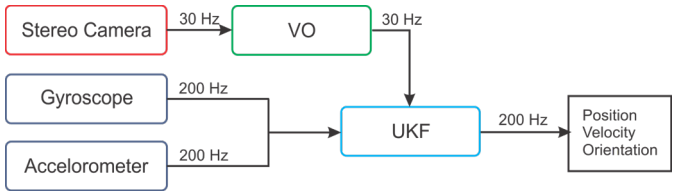


Fig. 7. The data flow diagram of the UKF used for fusing the visual odometry output with the IMU data

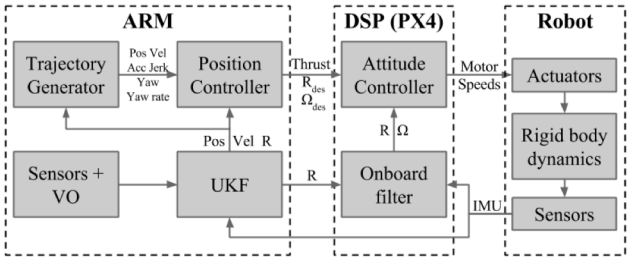


Fig. 8: An overview of the estimation and control pipeline running on the platform.

corridors that the robot should be inside in order to avoid collision with obstacles along the path. The generated trajectory is dynamically feasible and passes through the provided waypoints while keeping the robot in free space.

For the next step in miniaturization, we integrated the Snapdragon Flight board into the new Rotor-X Atom RX122 quadrotor vehicle [19] (Figure 9, 10) that became available during the second quarter of FY16. This is a palm-sized vehicle with a motor to motor span of 120 mm. As a COTS platform, the RotorX Atom comes with a Naze32 flight controller using a Cleanfight flight stack for manual RC control. In this configuration the vehicle weights 86 g without battery and about 160 g with the chosen battery. To achieve this integration, we built a custom frame element that reduces the foot print of the COTS Snapdragon Flight board from Intrinsyc by removing all mounting hardware and designing new custom made mounts. Since the RotorX Atom uses the same sensing and processing hardware as the Lumenier QAV180 platform, it has the same autonomous navigation capabilities as the QAV180.

We also integrated the autonomous landing capability reported in [1] into this system. In this configuration, the landing site detection pipeline itself runs at about 1 Hz on 1 ARM core, while visual-inertial state estimation uses 2 ARM cores. The Snapdragon processor includes automatic variation of clock rate to control temperature of the system. With the full state estimation and landing site detection system running, we logged temperature date with the rotors off and observed temperature rise from 45°C to a plateau between 80-85°C in about 30 seconds (Figure 11, top). Turning the rotors on brought the temperature down to 55°C in about 50 seconds (Figure 11, bottom). While this is a preliminary, qualitative test, it suggests that other temperature control alternatives are unnecessary, because rotor-induced airflow appears to produce adequate cooling. This is a data point for system integration work beyond the end of the MAST program.



Fig. 9. Palm-sized Rotor-X Atom RX122 quadrotor with integrated Snapdragon Flight board.

Regarding obstacle avoidance, the key issue is how to do obstacle detection. Within the MAST program, a high performance stereo vision ASIC was developed, but it currently exists in a board that is too heavy to add to what is already on the RotorX. Very near the end of the program, Qualcomm released object libraries for a stereo vision implementation that runs in the DSP on the Snapdragon board. With stereo cameras that could be added to the Snapdragon board, this would be a good option, but time and resources were no longer available to integrate and test this capability. Intel's Realsense active stereo vision system (i.e. with a patterned-light projector) could be an interesting alternative to consider in the future.

UPenn leveraged its close collaboration with Qualcomm to use a 250 gram quadrotor from Qualcomm that included the same Snapdragon processor module. UPenn built on their UKF-based visual-inertial state estimation work to provide onboard state updates at 500 Hz in order to support control of highly dynamic maneuvers, including paths around known obstacles and through small window-like structures [19]. This was extended to a motion planner that uses graph search with motion primitives to directly generate smooth polynomial trajectories. The onboard implementation was completed and benchmarked prior to the capstone, but in-flight tests were not. The onboard planning typically takes ~ 200-300 ms on 1 ARM core; polynomial trajectory evaluation at 200 Hz takes 20% of 1 ARM core.

B. Metrics for Performance Evaluation:

The goal of this project is to implement existing autonomous navigation and control capabilities on a true MAST-scale nano UAV platform. As such, the performance metric for this task is the successful integration and demonstration of individual components on the nano scale

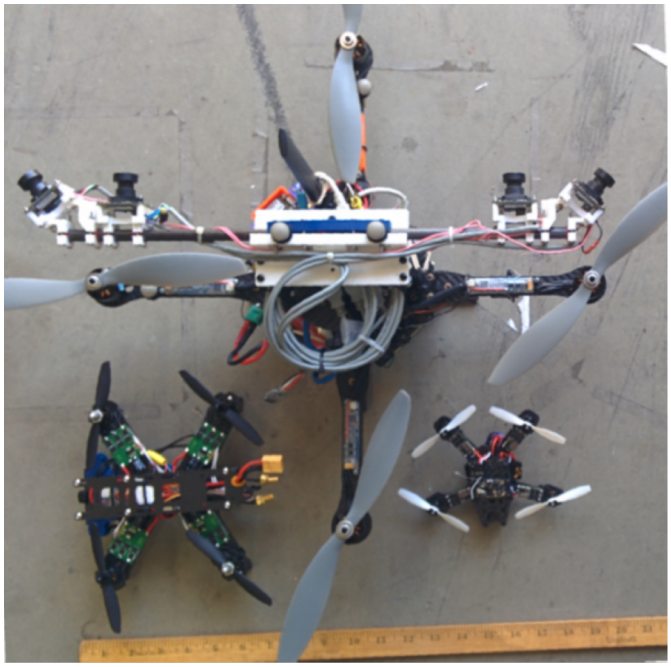


Fig. 10. Comparison of Rotor-X Atom (bottom right), Lumenier QAV180 (bottom left), and Asctec Pelican (top).

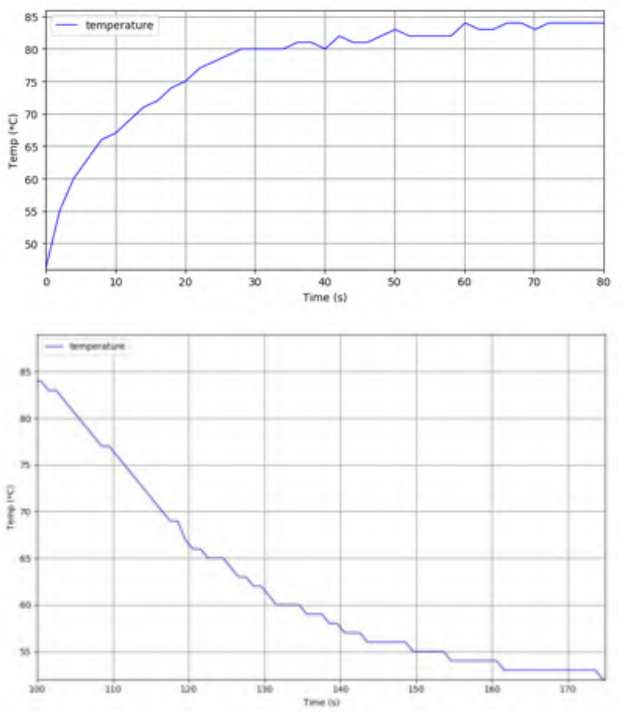


Fig. 11. Snapdragon Flight board temperature logs while running visual-inertial state estimation and visual landing site detection on the RotorX quadrotor. Top: with rotors off, the temperature rises to a plateau around 80-85°C in 30 seconds. Bottom: when rotors are turned on, temperature falls quickly.

platform, together with an evaluation that compares control stability and perception results with similar results on larger platforms that exist today within the MAST program. To measure control stability we will perform indoor tests in a defined environment and use ground truth from a VICON motion tracking system for error quantification. Perception system results will be compared by the quality of the achieved world representation (e.g. 3D maps), and the achievable computational throughput (frame rates, image resolutions, etc.). We aim for a full system that approaches palm-size, which would be an order of magnitude smaller than Ascending Technologies "Hummingbird" quadrotor which was the smallest platform capable of fully autonomous flight [1] at the start of this project.

V. Conclusions and Further Work

The expected end state of this task will be one of the smallest quadrotor vehicles available today that is capable of fully autonomous navigation in GPS-denied environments, including the execution of high level navigation tasks – as e.g. autonomous landing on high vantage points for surveillance applications. With this platform, we will showcase for the first time autonomous navigation technologies developed under the MAST program on a true MAST-scale vehicle.

References

[1] R. Brockers, M. Humenberger, S. Weiss, and L. Matthies, "Towards autonomous navigation of miniature UAV," IEEE Embedded Vision Workshop, held in association with the Computer Vision and Pattern Recognition Conference, 2014.

[2] Crazyflie Nano Quadcopter. https://www.bitcraze.io/crazyflie/

[3] L. Matthies, R. Brockers, Y. Kuwata, and S. Weiss. "Stereo vision-based obstacle avoidance for micro air vehicles using disparity space," IEEE International Conference on Robotics and Automation (ICRA), 2014.

[4] R. Brockers, S. Susca, D. Zhu, L. Matthies, "Fully self-contained vision-aided navigation and landing of a micro air vehicle independent from external sensor inputs," SPIE Symposium 8387: Unmanned Systems Technology XIV, 2012.

[5] V. Desaraju, N. Michael, M. Humenberger, R. Brockers, S. Weiss, L. Matthies, "Vision-based Landing Site Evaluation and Trajectory Generation Toward Rooftop Landing, " Proceedings of Robotics: Science and Systems, 2014.

[6] D. Mellinger, V. Kumar, "Minimum snap trajectory generation and control for quadrotors," Robotics and Automation (ICRA), 2011 IEEE International Conference on, 2520-2525, 2011.

[7] S. Weiss, M. W. Achtelik, S. Lynen, M. C. Achtelik, L. Kneip, M. Chli, and R. Siegwart, "Monocular Vision for Long-term Micro Aerial Vehicle State Estimation: A Compendium." J. Field Robotics, 30: 803–831. doi: 10.1002/rob.21466, 2013.

[8] S. Weiss, R. Brockers, L. Matthies, "4DoF drift free navigation using inertial cues and optical flow," Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on , vol., no., pp.4180,4186, 3-7 Nov. 2013.

[9] S. Weiss, R. Brockers, S. Albrektsen, L. Matthies. "Inertial Optical Flow for Throw-And-Go Micro Air Vehicles," IEEE Winter Conference on Applications of Computer Vision, 2015.

[10] http://qgroundcontrol.org/mavlink/start

[11] http://wiki.ros.org/mavros

[12] AeroVironment nano hummingbird. http://www.avinc.com/nano

[13] DelFly Explorer. http://www.delfly.nl/home.html.

[14] K. Ma, P. Chirarattanon, S. Fuller, and R.J. Wood. Controlled flight of a biologically inspired, insect-scale robot. *Science*, 340:603–607, 2013.

[15] AeroVironment Black Widow. https://www.avinc.com/uas/adc/black_widow/

[16] PD-100 Black Hornet PRS. www.proxdynamics.com/products/pd-100-blackhornet-prs

[17] Proto-X. www.estesrockets.com/helis- quads/quadcopters/004606-proto-x

[18] https://rotorxracing.com

[19] G. Loianno, C. Brunner, G. McGrath, V. Kumar, "Estimation, control, and planning for aggressive flight with a small quadrotor with a single camera and IMU," *IEEE RA-L Robotics and Automation Letters*, 2017. Also appeared in *ICRA 2017*

Metrics

Notes: *Please list metrics for the reporting quarter only and only Publications that are accepted and searchable on web.*

Students supported fully or partially (by name)

PhD Students: *Kartik Mohta (UPenn)*
Undergrad Students: *Norbert Nowak, Christian Stewart, Ryan D'Ambrosia (JPL)*

MAST Quarterly Metrics Summary Table (Numbers in this table should correspond to the list above):

MAST QUARTERLY REPORT METRICS TABLE (The data entered into this table should be for the reporting quarter only)	
Awards (Miscellaneous Recognitions and Awards)	0
Refereed/Journal papers	0
Conference presentations/proceedings	0
Joint Papers with ARL	0
Invention disclosures	0
Patent applications filed	0
Patents awarded	0
PhDs Award	0
Masters Awarded	0
Post Docs Supported	0
PhD Degree Students Supported	1
Master Degree Students Supported	0
Undergraduate Students Supported	3
CTA Member rotations into ARL	0
Visiting Researchers from ARL to the Member Sites	0
Employees hired by ARL from CTA students supported	0

6. MAST Program Bibliography

The following is the final bibliography for the MAST Program:

FY08

[1] Baras, J. S., and Hovareshti, P. (2008, December). Effects of topology in networked systems: Stochastic methods and small worlds. In *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on* (pp. 2973-2978). IEEE.

[2] Hovareshti, P., Baras, J. S., and Gupta, V. (2008, December). Average consensus over small world networks: A probabilistic framework. In *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on* (pp. 375-380). IEEE.

[3] Beekman, D. W., Mait, J. N., and Doligalski, T. L. (2008, July). Micro autonomous systems and technology at the army research laboratory. In *2008 IEEE National Aerospace and Electronics Conference* (pp. 159-162). IEEE.

FY09

[4] Ahmadzadeh, A., Motee, N., Jadbabaie, A., and Pappas, G. (2009, May). Multi-vehicle path planning in dynamically changing environments. In *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on* (pp. 2449-2454). IEEE.

[5] Akangah, S., Lingaiah, K., and Shivakumar (2009). Nylon-66 nanofiber interleaving on impact damage resistance of epoxy/carbon fiber composite laminate. In *50th AIAA/ASMC/ASCS/AHS/ASC Structures, Structural Dynamic, and Materials Conference*, Paul Springs, CA, 4-7 May 2009.

[6] Akangah, S., Lingaiah, K., and Shivakumar (2009, December). Nylon-66 nanofiber interleaving on impact damage resistance of epoxy/carbon fiber composite laminate. In *Energy Conversion Applications Power*.

[7] Akgul, M., Kim, B., Hung, L. W., Lin, Y., Li, W. C., Huang, W. L., and Nguyen, C. T. C. (2009, June). Oscillator far-from-carrier phase noise reduction via nano-scale gap tuning of micromechanical resonators. In *Transducers 2009-2009 International Solid-State Sensors, Actuators and Microsystems Conference* (pp. 798-801). IEEE.

[8] Aksaray, D., Dufresne, S., Raczynski, C., Zarovy, A., Cole, B., Mian, Z., and Mavris, D. (2009). MAST integrated design, analysis and simulation (MIDAS). In *Ankara International Aerospace Conference*.

[9] Antoulas, A. C. (2005). *Approximation of large-scale dynamical systems*. Society for Industrial and Applied Mathematics.

[10] Bachrach, A., He, R., and Roy, N. (2009). Autonomous flight in unknown indoor environments. *International Journal of Micro Air Vehicles*, *1*(4), 217-228.

[11] Baek, S. S., Ma, K. Y., and Fearing, R. S. (2009, October). Efficient resonant drive of flapping-wing robots. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 2854-2860). IEEE.

[12] Balakrisnan, B., Patil, S., and Smela, E. (2009). Patterning PDMS using a combination of wet and dry etching. *Journal of Micromechanics and Microengineering*, *19*(4), 047002.

[13] Baras, J. S., Hovareshti, P., and Perumal, S. (2009, August). Event triggered distributed collaborative control. In *Control Conference (ECC), 2009 European* (pp. 3797-3802). IEEE.

[14] Baras, J. S., Hovareshti, P., and Perumal, S. (2009, August). Event triggered distributed collaborative control. In *Control Conference (ECC), 2009 European* (pp. 3797-3802). IEEE.

[15] Beekman, D. W., Mait, J. N., and Doligalski, T. L. (2008, July). Micro autonomous systems and technology at the army research laboratory. In *2008 IEEE National Aerospace and Electronics Conference* (pp. 159-162). IEEE.

[16] Benedict, M., Ramasamy, M., Chopra, I., and Leishman, J. G. (2009, January). Experiments on the optimization of the MAV-scale cycloidal rotor characteristics towards improving their aerodynamic performance. In *Proceedings of the International Specialists Meeting on Unmanned Rotorcraft, Scottsdale, AZ* (pp. 20-22).

[17] Berman, S., Halász, Á., Hsieh, M. A., and Kumar, V. (2009). Optimized stochastic policies for task allocation in swarms of robots. *IEEE Transactions on Robotics*, *25*(4), 927-937.

[18] Billingsley, D., Slipher, G., Grauer, J., and Hubbard, J. (2009). Testing of a passively morphing ornithopter wing. *AIAA Paper*, *1828*.

[19] Birkmeyer, P., and Fearing, R. S. (2009, October). Dash: A resilient high-speed 15g hexapedal robot. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 418-419). IEEE.

[20] Bongsang Kim, B., Hung, L.-W., Akgul, M., Lin, Y., Li, W.-C., Huang, W.-L., Gurin, I., Borna, A. C., and Nguyen, T.-C. (2007). Engineering oscillator phase noise via nano-scale gap tuning of mems resonators. In *the 15th International Conference on Solid-State Sensors, Actuators and Microsystems*.

[21] Borno, R. T., Steinmeyer, J., and Maharbiz, M. M. (2008). Evaporation based electrical energy conversion in microfludic channels. In *26th Army Science Conference*, Orlando, Florida, December 1-4, 2008.

[22] Brooks, A., Makarenko, A., Kaupp, T., Durrant-Whyte, H., and Dellaert, F. (2009). Decentralised data fusion with dynamic topologies-A graphical model approach. *IFAC Proceedings Volumes*, *42*(20), 222-227.

[23] Cheng, P., Kumar, V., Arkin, R., Steinberg, M., and Hedrick, K. (2009). Cooperative control of multiple heterogeneous unmanned vehicles for coverage and surveillance [from the guest editors]. *IEEE Robotics & Automation Magazine*, *16*(2), 12-12.

[24] Choi, C., and Christensen, H. I. (2009, October). Cognitive vision for efficient scene processing and object categorization in highly cluttered environments. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on* (pp. 4267-4274). IEEE.

[25] Choi, C., Huckaby, J., III, J. R., Trevor, A. J. B., Case, J. P., and Christensen, H. I. (2009). Towards semantic perception for mobile manipulation.

[26] Choi, J., Oh, S., and Horowitz, R. (2009). Distributed learning and cooperative control for multi-agent systems. *Automatica*, *45*(12), 2802-2814.

[27] Faruque I., and Humbert J. S. (2009, June). Control models for dipertan flight dynamics, part i: longitudinal motion about hover, w/ minor revisions. *Journal of Theoretical Biology.*

[28] Finio, B. M., Eum, B., Oland, C., and Wood, R. J. (2009, October). Asymmetric flapping for a robotic fly using a hybrid power-control actuator. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 2755-2762). IEEE.

[29] Finio, B. M., Shang, J. K., and Wood, R. J. (2009, May). Body torque modulation for a microrobotic fly. In *Robotics and Automation. ICRA'09. IEEE International Conference on* (pp. 3449-3456). IEEE.

[30] Fink, J., Collins, T., Kumar, V., Mostofi, Y., Baras, J., and Sadler, B. (2009, May). A simulation environment for modeling and development of algorithms for ensembles of mobile microsystems. In *SPIE Defense, Security, and Sensing* (pp. 73180N-73180N). International Society for Optics and Photonics.

[31] Fink, J., Michael, N., Kim, S., and Kumar, V. (2011). Planning and control for cooperative manipulation and transportation with aerial robots. *The International Journal of Robotics Research*, *30*(3), 324-334.

[32] Ghaffarkhah, A., and Mostofi, Y. (2009). Communication-aware motion planning of robotic networks using navigation functions – centralized case. In *International Conference on Robot Communication and Coordination (RoboComm)*. IEEE.

[33] Ghaffarkhah, A., and Mostofi, Y. (2009, June). Communication-aware navigation functions for cooperative target tracking. In *American Control Conference, 2009. ACC'09.* (pp. 1316-1322). IEEE.

[34] Girard, A., and Pappas, G. J. (2006, December). Hierarchical control using approximate simulation relations. In *Decision and Control, 2006 45th IEEE Conference on* (pp. 264-269). IEEE.

[35] Goldman, D., Komsuoglu, H., and Koditschek, D. (2009). March of the Sandbots: legged robots are striding more successfully than ever with the help of some fleet-footed desert natives. *IEEE Spectrum*, *46*(4), 26-26.

[36] Grauer, J. A., and Hubbard, J. E. (2009). Inertial measurements from flight data of a flapping-wing ornithopter. *Journal of Guidance, Control, and Dynamics*, *32*(1), 326-331.

[37] Grauer, J., and Hubbard Jr, J. (2008, January). Development of a Sensor Suite for a Flapping-Wing UAV Platform. In *Proc. 46th AIAA Aerospace Sciences Meeting* (pp. 2008-224).

[38] Harmon, R., Grauer, J., Hubbard, J., Humbert, S., Roget, B., Siaraman, J., and Conroy, J. (2008). Experimental determination of ornithopter membrane wing shapes used for simple aerodynamic modeling. *AIAA Paper*, *6237*.

[39] Hong, W., and Sarabandi, K. (2008). *Low-profile vertically polarized omnidirectional antennas for ground sensor networks*. Michigan Univ., Ann Arbor.

[40] Hong, W., and Sarabandi, K. (2008, July). Miniaturized multi-element monopole antenna. In *2008 IEEE Antennas and Propagation Society International Symposium* (pp. 1-4). IEEE.

[41] Hong, W., and Sarabandi, K. (2009). Low-profile, multi-element, miniaturized monopole antenna. *IEEE Transactions on Antennas and Propagation*, *57*(1), 72-80.

[42] Hoover, A. M., and Fearing, R. S. (2009, October). Improvements to the stiffness properties of compliant mechanism designs with applications to mobile millirobots. In *IEEE Intelligent Robots and Systems Conference*, St. Louis, MO.

[43] Hoover, A. M., Steltz, E., and Fearing, R. S. (2008, September). RoACH: An autonomous 2.4 g crawling hexapod robot. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 26-33). IEEE.

[44] Hovareshti, P., and Baras, J. S. (2009, April). *Connectivity in networked systems: consensus problems and small world graphs*. ISR research day presentation, The Institute for Systems Research, University of Maryland, College Park, MD,

[45] Hovareshti, P., and Baras, J. S. (2009, April). *Connectivity in networked systems: consensus problems and small world graphs*. ISR research day presentation, The Institute for Systems Research, University of Maryland, College Park, MD,

[46] Humberta, J. S., Chopraa, I., Fearingb, R. S., Fullb, R. J., Woodc, R. J., and Dickinsond, M. H. (2009, May). Development of micromechanics for micro autonomous systems (ARL-MAST CTA program). In *SPIE Defense, Security, and Sensing* (pp. 73180L-73180L). International Society for Optics and Photonics.

[47] Jadbabaie, A., Ozdaglar, A., and Zargham, M. (2009, December). A distributed newton method for network optimization. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on* (pp. 2736-2741). IEEE.

[48] Jadbabaie, A., Ozdaglar, A., and Zargham, M. (2009, December). A distributed newton method for network optimization. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on* (pp. 2736-2741). IEEE.

[49] Jayaram, K., Mongeau, J. M., McRae, B., and Full, R. J. (2010, July). High-speed horizontal to vertical transitions in running cockroaches reveals a principle of robustness. *Integrative and Comparative Biology*, *50*, E83-E83. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[50] Jusufi, A., Gao, P. and Full, R.J., and Dudley, R. (2009). Gliding geckos actively use tails for turning. In *Society of Integrative and Comparative Biology Annual Meeting & Exhibition Final Program and Abstracts*. Boston, MA January 4-7, 2009.

[51] Kujawski, M., Pearse, J. D., and Smela, E. (2010). Elastomers filled with exfoliated graphite as compliant electrodes. *Carbon*, *48*(9), 2409-2417.

[52] Kuwata, Y., and Matthies, L. (2009, May). Robust landing using time-to-collision measurement with actuator saturation. In *SPIE Defense, Security, and Sensing* (pp. 73180Q-73180Q). International Society for Optics and Photonics.

[53] Kuwata, Y., and Matthies, L. (2009, May). Robust landing using time-to-collision measurement with actuator saturation. In *SPIE Defense, Security, and Sensing* (pp. 73180Q-73180Q). International Society for Optics and Photonics.

[54] Lakshminarayan, V. K., and Baeder, J. D. (2009). Computational investigation of small scale coaxial rotor aerodynamics in hover. *AIAA Paper*, *1069*, 5-9.

[55] Lakshminarayan, V. K., and Baeder, J. D. (2009). Numerical study of the effect of leading and trailing edge geometries on micro hovering rotor. In *AHS Specialists' Conference on Unmanned Rotorcraft*, Chandler, AZ, January 2009.

[56] Lakshminarayan, V. K., and Baeder, J. D. (2011). Computational investigation of microscale shrouded rotor aerodynamics in hover. *Journal of the American Helicopter Society*, *56*(4), 1-15.

[57] Langelaan, J. W., and Roy, N. (2009). Enabling new missions for robotic aircraft. *Science*, *326*(5960), 1642-1644.

[58] Li, C., Fearing, R., and Full, R. (2014, March). Insects traversing grass-like vertical compliant beams. *APS Meeting Abstracts*, *1*, 10012.

[59] Li, C., Umbanhowar, P. B., Komsuoglu, H., Koditschek, D. E., and Goldman, D. I. (2009). Sensitive dependence of the motion of a legged robot on granular media. *Proceedings of the National Academy of Sciences*, *106*(9), 3029-3034.

[60] Mait, J. N. (2009, May). The Army Research Laboratory's program on micro-autonomous systems and technology. In *SPIE Defense, Security, and Sensing* (pp. 73180K-73180K). International Society for Optics and Photonics.

[61] Makarenko, A., Brooks, A., Kaupp, T., Durrant-Whyte, H., and Dellaert, F. (2009, July). Decentralised data fusion: A graphical model approach. In *Information Fusion, 2009. FUSION'09. 12th International Conference on* (pp. 545-554). IEEE.

[62] Maladen, R. D., Ding, Y., Li, C., and Goldman, D. I. (2009). Undulatory swimming in sand: subsurface locomotion of the sandfish lizard. *Science*, *325*(5938), 314-318.

[63] Matei, I., Baras, J. S., and Jiang, T. (2009, July). A composite trust model and its application to collaborative distributed information fusion. In *Information Fusion, 2009. FUSION'09. 12th International Conference on* (pp. 1950-1957). IEEE.

[64] Matei, I., Martins, N. C., and Baras, J. S. (2009, June). Consensus problems with directed Markovian communication patterns. In *American Control Conference, 2009. ACC'09.* (pp. 1298-1303). IEEE.

[65] Matthey, L., Berman, S., and Kumar, V. (2009, May). Stochastic strategies for a swarm robotic assembly system. In *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on* (pp. 1953-1958). IEEE.

[66] Mayo, D. (2009). Universal power loading correlation for the efficiency evaluation of hover-capable devices. In *AIAA Mid-Atlantic Conference*.

[67] Mellinger, D., Kumar, V., and Yim, M. (2009, May). Control of locomotion with shape-changing wheels. In *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on* (pp. 1750-1755). IEEE.

[68] Michael, N., and Kumar, V. (2009). Planning and control of ensembles of robots with non-holonomic constraints. *The International Journal of Robotics Research*.

[69] Michael, N., Fink, J., and Kumar, V. (2008). Experimental testbed for large multirobot teams. *IEEE robotics & automation magazine*, *15*(1).

[70] Michael, N., Kim, S., Fink, J., and Kumar, V. (2009, January). Kinematics and statics of cooperative multi-robot aerial manipulation with cables. In *ASME 2009 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (pp. 83-91). American Society of Mechanical Engineers.

[71] Moallem, M., East, J., and Sarabandi, K. (2009, June). Optimally designed membrane-supported grounded CPW structure for submillimeter-wave applications. In *2009 IEEE Antennas and Propagation Society International Symposium* (pp. 1-4). IEEE.

[72] Mongeau, J. M., Jayaram, K., Lee, J., Full, R. J., and Cowan, N. (2010, July). Mechanical feedback of antenna-substrate interaction simplifies cockroach antennal navigation. *Integrative and Comparative Biology*, *50*, E120-E120. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[73] Moore, T., Burden, S., Revzen, S., and Full, R. J. (2010). Adding inertia to test stability predictions in rapid running insects. In *Society of Integrative and Comparative Biology Annual Meeting & Exhibition Final Program and Abstracts*. Seattle, WA. January 3-7, 2010.

[74] Mostofi, Y. (2008, May). Communication-aware motion planning in fading environments. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on* (pp. 3169-3174). IEEE.

[75] Mostofi, Y., and Murray, R. M. (2009). Kalman filtering over wireless fading channels-How to handle packet drop. *International Journal of Robust and Nonlinear Control*, *19*(18), 1993-2015.

[76] Mostofi, Y., and Murray, R. M. (2009). To drop or not to drop: design principles for Kalman filtering over wireless fading channels. *IEEE Transactions on Automatic Control*, *54*(2), 376-381.

[77] Mostofi, Y., Gonzalez-Ruiz, A., Gaffarkhah, A., and Li, D. (2009, October). Characterization and modeling of wireless channels for networked robotic and control systems-a comprehensive overview. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on* (pp. 4849-4854). IEEE.

[78] Mostofi, Y., Gonzalez-Ruiz, A., Gaffarkhah, A., and Li, D. (2009, October). Characterization and modeling of wireless channels for networked robotic and control systems-a comprehensive overview. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on* (pp. 4849-4854). IEEE.

[79] Oh, J., Thiel, M., Hong, W., and Sarabandi, K. (2009, June). Indoor wave propagation measurements and modeling for evaluation of coverage enhancement using a repeater system. In *2009 IEEE Antennas and Propagation Society International Symposium*.

[80] Ramsey, J. (2009). Flow visualization and particle image velocimetry of an insect-based flapping wing mechanism. In *AIAA Mid-Atlantic Student Conference*.

[81] Revzen, S., and Full, R. J. (2009). From neuromechanics to robot controllers: Bio-inspired design. *Comparative Biochemistry and Physiology Part A: Molecular & Integrative Physiology*, *153*(2), S136.

[82] Rogers, J. G., and Christensen, H. I. (2009, October). Normalized graph cuts for visual slam. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on* (pp. 918-923). IEEE.

[83] Rogers, J. G., and Christensen, H. I. (2009, October). Normalized graph cuts for visual slam. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on* (pp. 918-923). IEEE.

[84] Sadeghi, M., Kim, H., and Najafi, K. (2010, January). Electrostatically driven micro-hydraulic actuator arrays. In *Micro Electro Mechanical Systems (MEMS), 2010 IEEE 23rd International Conference on* (pp. 15-18). IEEE.

[85] Schwager, M., Julian, B. J., and Rus, D. (2009, May). Optimal coverage for multiple hovering robots with downward facing cameras. In *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on* (pp. 3515-3522). IEEE.

[86] Schwager, M., Julian, B. J., and Rus, D. (2009, May). Optimal coverage for multiple hovering robots with downward facing cameras. In *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on* (pp. 3515-3522). IEEE.

[87] Shivakumar, K., Lingaiah, S., Chen, H., Akangah, P., Swaminathan, G., and Russell, L. (2009). Polymer nanofabric interleaved composite laminates. *AIAA journal*, *47*(7), 1723-1729.

[88] Singh, B., and Chopra, I. (2008). Insect-based hover-capable flapping wings for micro air vehicles: experiments and analysis. *AIAA journal*, *46*(9), 2115-2135.

[89] Sitaraman, J., Roget, B., Harmon, R., Grauer, J., Conroy, J., Hubbard, J., and Humbert, S. (2008). A computational study of flexible wing ornithopter flight. In *26th AIAA Applied Aerodynamics Conference* (p. 6397).

[90] Slaton, Cohen, Ding, Umbanhowar, Goldman, et al. (2009). Integrating a Hierarchy of Simulation Tools. In *IROS 2009*.

[91] Spence, A. J., Seipel, J., Revzen, S., Mullens, C., Yeats, K., and Full, R. J. (2009). Insects running on elastic surfaces: The role of feedforward control. *Comparative Biochemistry and Physiology Part A: Molecular & Integrative Physiology*, *153*(2), S137.

[92] St. Clair, J., Jackson, G. S., and Cadou, C. P. (2009). Scaling of fuel cell systems for micro-power applications. In *9th International Workshop on Micro and Nanotechnology for Power Generation*.

[93] Toshev, A., Makadia, A., and Daniilidis, K. (2009). Shape-based detection of moving objects in videos. In *IEEE Comp. Vision Pattern Recognition*.

[94] Toshev, A., Makadia, A., and Daniilidis, K. (2009). Shape-based detection of moving objects in videos. In *IEEE Comp. Vision Pattern Recognition*.

[95] Umbanhowar, Ding, Li, and Goldman. (2009). Volume fraction variation drives crossover in granular collision dynamics. *Physical Review Letters*.

[96] Wu, J., Christensen, H. I., and Rehg, J. M. (2009, October). Visual place categorization: Problem, dataset, and algorithm. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on* (pp. 4763-4770). IEEE.

[97] Wu, J., Christensen, H. I., and Rehg, J. M. (2009, October). Visual place categorization: Problem, dataset, and algorithm. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on* (pp. 4763-4770). IEEE.

[98] Yamaleev, N. K., and Carpenter, M. H. (2009). A systematic methodology for constructing high-order energy stable WENO schemes. *Journal of Computational Physics*, *228*(11), 4248-4272.

[99] Yamaleev, N., and Carpenter, M. H. (2009, January). High energy stable WENO schemes. In *AIAA*, Orlando, FL. Paper 2009-1135.

[100] Yang, K., El-Emawy, M. A., Gu, T. Y., Stintz, A., and Lester, L. F. (2009, October). GaAs-based InAs/InGaAs quantum dot solar cells for concentration applications. In *Frontiers in Optics* (p. FMB4). Optical Society of America.

[101] Zhu, X., Lee, V., Phillips, J., and Mortazawi, A. (2009). An intrinsically switchable FBAR filter based on barium titanate thin films. *IEEE Microwave and Wireless Components Letters*, *19*(6), 359-361.

[102] Zhu, X., Lee, V., Phillips, J., and Mortazawi, A. (2009, June). Intrinsically switchable contour mode acoustic wave resonators based on

barium titanate thin films. In *Microwave Symposium Digest, 2009. MTT'09. IEEE MTT-S International* (pp. 93-96). IEEE.

FY10

[103] Akangah, P., Lingaiah, S., and Shivakumar, K. (2010). Effect of nylon-66 nano-fiber interleaving on impact damage resistance of epoxy/carbon fiber composite laminates. *Composite Structures*, *92*(6), 1432-1439.

[104] Albright, J., and Sirohi, J. (2010). Conceptual design of a hover-capable high-speed multi-role MAV. In *66th Annual Forum of the AHS*, Phoenix, AZ, May 11-13, 2010.

[105] Anati, R., and Daniilidis, K. (2009). Constructing topological maps using markov random fields and loop-closure detection. In *Advances in Neural Information Processing Systems* (pp. 37-45).

[106] Ayanian, N., and Kumar, V. (2010, May). Abstractions and controllers for groups of robots in environments with obstacles. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on* (pp. 3537-3542). IEEE.

[107] Ayanian, N., Kumar, V., and Koditschek, D. (2011). Synthesis of controllers to create, maintain, and reconfigure robot formations with communication constraints. In *Robotics Research* (pp. 625-642). Springer Berlin Heidelberg.

[108] Baras, J. S., and Hovareshti, P. (2009, December). Efficient and robust communication topologies for distributed decision making in networked systems. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on* (pp. 3751-3756). IEEE.

[109] Benedict, M., Chopra, I., Mattaboni, M., and Masarati, P. (2010, April). Aeroelastic analysis of a MAV-scale cycloidal rotor. In *51st SDM Conference*, Orlando, FL.

[110] Benedict, M., Chopra, I., Ramasamy, M., and Leishman, J. G. (2008). Experimental investigation of the cycloidal-rotor concept for a hovering micro air vehicle. *Annual Forum Proceedings-American Helicopter Society*, *64*(2), 1071. American Helicopter Society, Inc.

[111] Benedict, M., Jarugumilli, T., and Chopra, I. (2010). Experimental performance optimization of a MAV-scale cycloidal rotor. In *Proceedings of the AHS Specialists' Meeting on Aeromechanics* (pp. 119-135). Curran Associates, Inc.

[112] Benedict, M., Ramasamy, M., and Chopra, I. (2010). Improving the aerodynamic performance of micro-air-vehicle-scale cycloidal rotor: an experimental approach. *Journal of aircraft*, *47*(4), 1117-1125.

[113] Benedict, M., Ramasamy, M., Chopra, I., and Leishman, J. G. (2010). Performance of a cycloidal rotor concept for micro air vehicle applications. *Journal of the American Helicopter Society*, *55*(2), 22002-22002.

[114] Bhattacharya, S., Michael, N., and Kumar, V. (2013). Distributed coverage and exploration in unknown non-convex environments. In *Distributed autonomous robotic systems* (pp. 61-75). Springer Berlin Heidelberg.

[115] Billingsley, D., Slipher, G., Grauer, J., and Hubbard, J. (2009). Testing of a passively morphing ornithopter wing. *AIAA Paper*, *1828*.

[116] Birkmeyer, P., Peterson, K., and Fearing, R. S. (2009, October). DASH: A dynamic 16g hexapedal robot. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 2683-2689). IEEE.

[117] Borno, R. T., Steinmeyer, J. D., and Maharbiz, M. M. (2009). Charge-pumping in a synthetic leaf for harvesting energy from evaporation-driven flows. *Applied Physics Letters*, *95*(1), 013705.

[118] Brooks, A., Makarenko, A., Kaupp, T., Durrant-Whyte, H., and Dellaert, F. (2009). Decentralised Data Fusion with Dynamic Topologies-A Graphical Model Approach, *IFAC Proceedings Volumes*, *42*(20), 222-227.

[119] Chen, H., Hovareshti, P., and Baras, J. S. (2010, June). Distributed collaborative controlled autonomous vehicle systems over wireless networks. In *Control & Automation (MED), 2010 18th Mediterranean Conference on* (pp. 1695-1700). IEEE.

[120] Cunningham, A., Paluri, B., and Dellaert, F. (2010). Fully distributed SLAM using constrained factor graphs. In *Proceedings of the IEEE Conference on Intelligent Robots and Systems*.

[121] Dellaert, F., Carlson, J., Ila, V., Ni, K., and Thorpe, C. E. (2010, October). Subgraph-preconditioned conjugate gradients for large scale SLAM. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on* (pp. 2566-2571). IEEE.

[122] Derenick, J., Kumar, V., and Jadbabaie, A. (2010, May). Towards simplicial coverage repair for mobile robot teams. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on* (pp. 5472-5477). IEEE.

[123] Derenick, J., Spletzer, J., and Kumar, V. (2010, December). A semidefinite programming framework for controlling multi-robot systems in dynamic environments. In *Decision and Control (CDC), 2010 49th IEEE Conference on* (pp. 7172-7177). IEEE.

[124] Dickson, W. B, Polidoro, P. J., Tanner, M. M., Dickinson, M. H. (2010). A linear systems analysis of the yaw dynamics of the fruit fly Drosophila. J. Exp. Biol., in press.

[125] Eun, C. K., and Gianchandani, Y. B. (2010). A wireless-enabled radiation detector using micromachined steel and glass elements in a TO-5 package. *Stainless steel*, *302*, 304.

[126] Eun, C. K., and Gianchandani, Y. B. (2011, January). A microdischarge-based radiation detector utilizing stacked electrode arrays in a TO-5 package. In *Micro Electro Mechanical Systems (MEMS), 2011 IEEE 24th International Conference on* (pp. 668-671). IEEE.

[127] Eun, C. K., Gharpurey, R., and Gianchandani, Y. B. (2010). Wireless signaling of beta detection using microdischarges. *Journal of Microelectromechanical Systems*, *19*(4), 785-793.

[128] Ferriss, M., Lin, D. T., and Flynn, M. P. (2009, September). A fractional-N PLL modulator with flexible direct digital phase modulation. In *2009 IEEE Custom Integrated Circuits Conference* (pp. 49-52). IEEE.

[129] Finio, B. M., and Wood, R. J. (2010). Distributed power and control actuation in the thoracic mechanics of a robotic insect. *Bioinspiration & biomimetics*, *5*(4), 045006.

[130] Finio, B. M., and Wood, R. J. (2011, September). Optimal energy density piezoelectric twisting actuators. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 384-389). IEEE.

[131] Finio, B. M., Whitney, J. P., and Wood, R. J. (2010, October). Stroke plane deviation for a microrobotic fly. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on* (pp. 3378-3385). IEEE.

[132] Fisher, T. C., Carpenter, M. H., Yamaleev, N. K., and Frankel, S. H. (2011). Boundary closures for fourth-order energy stable weighted essentially non-oscillatory finite-difference schemes. *Journal of Computational Physics*, *230*(10), 3727-3752.

[133] Ghaffarkhah, A., and Mostofi, Y. (2010, June). Channel learning and communication-aware motion planning in mobile networks. In *American Control Conference (ACC), 2010* (pp. 5413-5420). IEEE.

[134] Gil, S., Schwager, M., Julian, B., and Rus, D. (2010). Distributed control for optimization of a communication network. In *Proceedings of the 2010 IEEE Int'l Conference on Robotics and Automation*, Anchorage, AK.

[135] Godbehere, A. & Sastry, S. (2010). Polar coordinate methods for the design and analysis of switched linear systems. In the *Proceedings of HSCC 2010, the International Conference on Hybrid Systems: Computation and Control*.

[136] Goldman, D. I., and Hu, D. L. (2010). Wiggling Through the World: The mechanics of slithering locomotion depend on the surroundings. *American Scientist*, *98*(4), 314-323.

[137] Gonçalves, V. M., Pimenta, L. C., Maia, C. A., Pereira, G. A., Dutra, B. C., Michael, N., ... & Kumar, V. (2010, May). Circulation of curves using vector fields: actual robot experiments in 2D and 3D workspaces. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on* (pp. 1136-1141). IEEE.

[138] Gonzalez, H., Vasudevan, R., Kamgarpour, M., Sastry, S. S., Bajcsy, R., and Tomlin, C. J. (2010, April). A descent algorithm for the optimal control of constrained nonlinear switched dynamical systems. In *Proceedings of the 13th ACM international conference on Hybrid systems: computation and control* (pp. 51-60). ACM.

[139] Grauer, J. A., and Hubbard, J. E. (2009). Inertial measurements from flight data of a flapping-wing ornithopter. *Journal of Guidance, Control, and Dynamics*, *32*(1), 326-331.

[140] Grauer, J. A., and Hubbard, J. E. (2009). Multibody model of an ornithopter. *Journal of Guidance, Control, and Dynamics*, *32*(5), 1675-1679.

[141] Grauer, J., and Hubbard Jr, J. (2008, January). Development of a sensor suite for a flapping-wing UAV platform. In *Proc. 46th AIAA Aerospace Sciences Meeting* (pp. 2008-224).

[142] Grauer, J., and Hubbard, J. (2009). Identification and integration of ornithopter actuator models. In *AIAA Atmospheric Flight Mechanics Conference*, Chicago, IL, AIAA Paper (Vol. 5937).

[143] Grauer, J., Ulrich, E., Hubbard, J., Pines, D., and Humbert, J. (2010, January). Model structure determination of an ornithopter aerodynamics model from flight data. Orlando, FL.

[144] Gu, T., El-Emawy, M. A., Yang, K., Stintz, A., and Lester, L. F. (2009). Resistance to edge recombination in GaAs-based dots-in-a-well solar cells. *Applied Physics Letters*, *95*(26), 261106.

[145] Harmon, R., Grauer, J., Hubbard, J., Humbert, S., Roget, B., Siaraman, J., and Conroy, J. (2008). Experimental determination of ornithopter membrane wing shapes used for simple aerodynamic modeling. *AIAA Paper*, *6237*.

[146] He, R., Bachrach, A., Achtelik, M., Geramifard, A., Gurdan, D., Prentice, S., and Roy, N. (2010). On the design and use of a micro air vehicle to track and avoid adversaries. *The International Journal of Robotics Research*, *29*(5), 529-546.

[147] Hoover, A. M., Burden, S., Fu, X. Y., Sastry, S. S., and Fearing, R. S. (2010, September). Bio-inspired design and dynamic maneuverability of a minimally actuated six-legged robot. In *Biomedical Robotics and*

Biomechatronics (BioRob), 2010 3rd IEEE RAS and EMBS International Conference on (pp. 869-876). IEEE.

[148] Kaess, M., and Dellaert, F. (2009). Covariance recovery from a square root information matrix for data association. *Robotics and autonomous systems*, *57*(12), 1198-1210.

[149] Kira, Z. (2009, November). Mapping grounded object properties across perceptually heterogeneous embodiments. In *FLAIRS Conference* (p. 37).

[150] Kira, Z. (2009, October). Transferring embodied concepts between perceptually heterogeneous robots. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on* (pp. 4650-4656). IEEE.

[151] Kira, Z. (2010, May). Inter-robot transfer learning for perceptual classification. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1* (pp. 13-20). International Foundation for Autonomous Agents and Multiagent Systems.

[152] Kira, Z., Arkin, R. C., and Collins, T. R. (2010, April). A design process for robot capabilities and missions applied to micro-autonomous platforms. In *SPIE Defense, Security, and Sensing* (pp. 767911-767911). International Society for Optics and Photonics.

[153] Kujawski, M., Pearse, J. D., and Smela, E. (2010). Elastomers filled with exfoliated graphite as compliant electrodes. *Carbon*, *48*(9), 2409-2417.

[154] Kujawski, M., Pearse, J., and Smela, E. (2010, March). PDMS/graphite stretchable electrodes for dielectric elastomer actuators. In *SPIE Smart Structures and Materials+ Nondestructive Evaluation and Health Monitoring* (pp. 76420R-76420R). International Society for Optics and Photonics.

[155] Kumar, M., Garg, D. P., and Kumar, V. (2010). Segregation of heterogeneous units in a swarm of robotic agents. *IEEE Transactions on Automatic Control*, *55*(3), 743-748.

[156] Lakshminarayan, V. K., and Baeder, J. D. (2011). Computational investigation of microscale shrouded rotor aerodynamics in hover. *Journal of the American Helicopter Society*, *56*(4), 1-15.

[157] Lee, V., Sis, S. A., Zhu, X., and Mortazawi, A. (2010). Interdigitated contour mode resonators based on barium titanate thin films. In *IEEE International Microwave Symposium Digest*.

[158] Lee, V., Sis, S. A., Zhu, X., and Mortazawi, A. (2010, May). Intrinsically switchable interdigitated barium titanate thin film contour mode resonators. In *Microwave Symposium Digest (MTT), 2010 IEEE MTT-S International* (pp. 1448-1450). IEEE.

[159] Lee, V., Sis, S. A., Zhu, X., and Mortazawi, A. (2010, September). Intrinsically switchable ferroelectric BAW resonators and filters. In *Microwave Conference (EuMC), 2010 European* (pp. 803-806). IEEE.

[160] Lee, V., Sis, S. A., Zhu, X., and Mortazawi, A. (2010, September). Intrinsically switchable ferroelectric BAW resonators and filters. In *Microwave Conference (EuMC), 2010 European* (pp. 803-806). IEEE.

[161] Lee, V.; Sis, S. A.; Mortazawi, A. (2010). Ferroelectric resonators and filters for reconfigurable radios. In *Army Science Conference*.

[162] Lee, V.; Sis, S. A.; Zhu, X.; Mortazawi, A. (2010, December). Interdigitated contour mode resonators based on barium titanate thin films. In *IEEE International Microwave Symposium Digest*.

[163] Lester, L. F., Yang, K., El-Emawy, M. A., Saiz, T., Clark, R. A., Hollowell, A., and Lavrova, O. (2010, April). Flexible solar cells for micro-autonomous systems technology. In *SPIE Defense, Security, and Sensing* (pp. 76790Y-76790Y). International Society for Optics and Photonics.

[164] Lester, L., Yang, K., El-Emawy, M., Saiz, T., Clark, R. A., Hollowell, A., and Lavrova, O. (2010, March). Solar power for unmanned autonomous vehicles. *SPIE newsroom article*, #10.1117/2.1201001.002601.

[165] Lester, L., Yang, K., El-Emawy, M., Saiz, T., Clark, R., Hollowell, A., and Lavrova, O. Solar power for unmanned autonomous vehicles.

[166] Li, C., Umbanhowar, P. B., Komsuoglu, H., and Goldman, D. I. (2010). The effect of limb kinematics on the speed of a legged robot on granular media. *Experimental mechanics*, *50*(9), 1383-1393.

[167] Lin, D. T., Li, L., Farahani, S., and Flynn, M. P. (2012). A flexible 500 MHz to 3.6 GHz wireless receiver with configurable DT FIR and IIR filter embedded in a 7b 21 MS/s SAR ADC. *IEEE Transactions on Circuits and Systems I: Regular Papers*, *59*(12), 2846-2857.

[168] Maladen, R. D., Ding, Y., Umbanhowar, P. B., Kamor, A., and Goldman, D. I. (2010). Biophysically inspired development of a sand-swimming robot. In *Robotics: Science & Systems (RSS)*.

[169] Masarati, P., Quaranta, G., Gennaretti, M., and Serafini, J. (2010, May). Aeroservoelastic analysis of rotorcraft-pilot coupling: a parametric study. In *Proceedings of the American Helicopter Society 66th Annual Forum*.

[170] Mattaboni, M., Benedict, M., Masarati, P., and Chopra, I. (2010). MAV-scale cycloidal rotor multibody aeroelastic analysis. In *1st Joint International Conference on Multibody System Dynamics (IMSD 2010)*, Lappeenranta, Finland, May 25-27, 2010.

[171] Mazouchova, N., Gravish, N., Savu, A., and Goldman, D. I. (2010). Limb-ground interaction during terrestrial locomotion of hatchling sea

turtles. in *Biology Letters*, special invited issue on Control and Dynamics of Animal Movement.

[172] Mazouchova, N., Gravish, N., Savu, A., and Goldman, D. I. (2010). Utilization of granular solidification during terrestrial locomotion of hatchling sea turtles. *Biology letters*, *6*(3), 398-401.

[173] Mellinger, D., and Kumar, V. (2010, May). Control and planning for vehicles with uncertainty in dynamics. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on* (pp. 960-965). IEEE.

[174] Michael, N., and Kumar, V. (2011). Control of ensembles of aerial robots. *Proceedings of the IEEE*, *99*(9), 1587-1602.

[175] Michael, N., Mellinger, D., Lindsey, Q., and Kumar, V. (2010). The grasp multiple micro-uav testbed. *IEEE Robotics & Automation Magazine*, *17*(3), 56-65.

[176] Motee, N., Jadbabaie, A., and Pappas, G. (2010, June). Path planning for multiple robots: An alternative duality approach. In *American Control Conference (ACC), 2010* (pp. 1611-1616). IEEE.

[177] Motee, N., Jadbabaie, A., and Pappas, G. (2010, May). A duality approach to path planning for multiple robots. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on* (pp. 935-940). IEEE.

[178] Nieto-Granda, C., Rogers, J. G., Trevor, A. J., and Christensen, H. I. (2010, October). Semantic map partitioning in indoor environments using regional analysis. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on* (pp. 1451-1456). IEEE.

No papers

[179] Oh, J., and Sarabandi, K. (2010). Low-profile miniaturized antenna with high efficiency for ground sensor networks. In *the 27th Army Science Conference*.

[180] Oh, J., and Sarabandi, K. (2010, July). A novel high gain low profile miniaturized vertically polarized antenna. In *2010 IEEE Antennas and Propagation Society International Symposium* (pp. 1-4). IEEE.

[181] Oh, J., and Sarabandi, K. Low-profile miniaturized antenna with high efficiency for ground sensor networks. *Ann Arbor*, *1001*, 48109-2122.

[182] Oh, S., and Choi, J. (2010, December). Distributed learning in mobile sensor networks using cross validation. In *Decision and Control (CDC), 2010 49th IEEE Conference on* (pp. 3845-3850). IEEE.

[183] Oh, S., Xu, Y., and Choi, J. (2010, December). Sparse Gaussian processes for mobile sensor networks. In *CDC*.

[184] Panagou, D., Tanner, H. G., and Kyriakopoulos, K. J. (2010, May). Dipole-like fields for stabilization of systems with Pfaffian constraints. In

Robotics and Automation (ICRA), 2010 IEEE International Conference on (pp. 4499-4504). IEEE.

[185] Park, Y., and Wentzloff, D. D. (2010, May). A cyclic vernier time-to-digital converter synthesized from a 65nm CMOS standard library. In *Proceedings of 2010 IEEE International Symposium on Circuits and Systems* (pp. 3561-3564). IEEE.

[186] Pimenta, L. C., Schwager, M., Lindsey, Q., Kumar, V., Rus, D., Mesquita, R. C., and Pereira, G. A. (2009). Simultaneous coverage and tracking (SCAT) of moving targets with robot networks. In *Algorithmic foundation of robotics VIII* (pp. 85-99). Springer Berlin Heidelberg.

[187] Piovesan, J. L., and Tanner, H. G. (2009, May). Randomized model predictive control for robot navigation. In *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on* (pp. 94-99). IEEE.

[188] Rogers, J. G., Trevor, A. J., Nieto-Granda, C., and Christensen, H. I. (2010, October). Slam with expectation maximization for moveable object tracking. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on* (pp. 2077-2082). IEEE.

[189] Roget, B., Sitaraman, J., Harmon, R., Grauer, J., Hubbard, J., and Humbert, S. (2009). Computational study of flexible wing ornithopter flight. *Journal of Aircraft*, *46*(6), 2016-2031.

[190] Sadeghi, M., Peterson, R. L., Peterson, K., Fearing, R., and Najafi, K. (2010, December). Air-flow sensing on autonomous mobile platforms using micro-scale hot-wire anemometry. In *Proc. 27th Army Science Conference*.

[191] Sarabandi, K., and Song, Y. J. (2011). Subwavelength radio repeater system utilizing miniaturized antennas and metamaterial channel isolator. *IEEE Transactions on Antennas and Propagation*, *59*(7), 2683-2690.

[192] Seshadri, P., Benedict, M., and Chopra, I. (2009). Understanding insect-based flapping flight from a micro air vehicle perspective. In *Proceedings of the American Helicopter Society International Specialists Meeting on Unmanned Rotorcraft*, Phoenix, AZ January 19-22, 2009.

[193] Seshadri, P., Benedict, M., and Chopra, I. (2010). Experimental investigation of an avian-based flapping wing concept for a micro air vehicle. In *Proceedings of the American Helicopter Society Aeromechanics Conference*, San Francisco, CA, January 20-22, 2010.

[194] Shiu, K. T., Zimmerman, J., Wang, H., and Forrest, S. R. (2009). Ultrathin film, high specific power InP solar cells on flexible plastic substrates. *Applied Physics Letters*, *95*(22), 223503.

[195] Shiu, K.-T., Zimmerman, J., Wang, H., and Forrest, S. R. (2009). Thin-film ITO/InP solar cells on flexible plastic substrates. In *Materials Research Society Fall Meeting*, Boston. MA.

[196] Song, Y. J., and Sarabandi, K. (2010, July). Small radio repeater system for enhancement of wireless connectivity. In *2010 IEEE Antennas and Propagation Society International Symposium* (pp. 1-4). IEEE.

[197] Song, Y. J., and Sarabandi, K. (2010, July). Small radio repeater system for enhancement of wireless connectivity. In *2010 IEEE Antennas and Propagation Society International Symposium* (pp. 1-4). IEEE.

[198] Tanner, H. G., and Piovesan, J. L. (2010). Randomized receding horizon navigation. *IEEE Transactions on Automatic Control*, *55*(11), 2640-2644.

[199] Trevor, A. J., Rogers, J. G., Nieto, C., and Christensen, H. I. (2010, May). Applying domain knowledge to SLAM using virtual measurements. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on* (pp. 5389-5394). IEEE.

[200] Ulam, P. D., Kira, Z., Arkin, R. C., and Collins, T. R. (2010, April). Mission specification and control for unmanned aerial and ground vehicles for indoor target discovery and tracking. In *SPIE Defense, Security, and Sensing* (pp. 769414-769414). International Society for Optics and Photonics.

[201] Umbanhowar, P., and Goldman, D. I. (2010). Granular impact and the critical packing state. *Physical Review E*, *82*(1), 010301.

[202] Watteyne, T., Lanzisera, S., Mehta, A., and Pister, K. S. (2010, May). Mitigating multipath fading through channel hopping in wireless sensor networks. In *Communications (ICC), 2010 IEEE International Conference on* (pp. 1-5). IEEE.

[203] Watteyne, T., Mehta, A., and Pister, K. (2009, October). Reliability through frequency diversity: why channel hopping makes sense. In *Proceedings of the 6th ACM symposium on Performance evaluation of wireless ad hoc, sensor, and ubiquitous networks* (pp. 116-123). ACM.

[204] Wise, K. D. (2010). Fantastic voyage: from integrated sensors to microsystems. Invited keynote in *MEMS Technology Summit: Lessons from the Past and Visions of the Future*, Stanford University, October 2010.

[205] Wise, K. D. (2010). Microelectronics in the "more than Moore" era. In *68th Device Research Conference*.

[206] Yamaleev, N. K., and Carpenter, M. H. (2011). Energy stable WENO schemes of arbitrary order. In *Computational Fluid Dynamics 2010* (pp. 81-87). Springer Berlin Heidelberg.

[207] Yang, K., El-Emawy, M. A., Gu, T. Y., Stintz, A., and Lester, L. F. (2009, October). GaAs-based InAs/InGaAs quantum dot solar cells for concentration applications. In *Frontiers in Optics* (p. FMB4). Optical Society of America.

[208] Yang, K., Lakshminarayan, V. K., and Baeder, J. D. (2010, May). Simulation of a cycloidal rotor system using an overset RANS solver. In *at American Helicopter Society 66th Annual Forum Proceedings, Phoenix, AZ*.

FY11

[209] Bachrach, A., Prentice, S., He, R., and Roy, N. (2011). Range–robust autonomous navigation in GPS-denied environments. *Journal of Field Robotics*, *28*(5), 644-666.

[210] Baras, J. S., and Hovareshti, P. (2010, July). Efficient communication infrastructures for distributed control and decision making in networked stochastic systems. In *Proceedings of the 19th International Symposium on Mathematical Theory of Networks and Systems (MTNS 2010)* (pp. 5-9).

[211] Beach, K. T. M., Reidy, S. M., Gordenker, R. J. M., and Wise, K. D. (2011, January). A low-mass high-speed μGC separation column with built-in fluidic chip-to-chip interconnects. In *Micro Electro Mechanical Systems (MEMS), 2011 IEEE 24th International Conference on* (pp. 813-816). IEEE.

[212] Benedict, M., Chopra, I., Mattaboni, M., and Masarati, P. (2010, April). Aeroelastic analysis of a MAV-scale cycloidal rotor. In *51st SDM Conference*, Orlando, FL.

[213] Benedict, M., Gupta, R., and Chopra, I. (2011). Design, development and flight testing of a twin-rotor cyclocopter micro air vehicle. In *Proceedings of the 67th Annual National Forum of the American Helicopter Society*, Virginia Beach, VA, May 3–5, 2011.

[214] Benedict, M., Jarugumilli, T., and Chopra, I. (2009). Design and development of a hover-capable cyclocopter MAV. *Proceedings of the 65th Annual National Froum of the American Helicopter Society. Grapevine: Curran Associates*.

[215] Benedict, M., Jarugumilli, T., and Chopra, I. (2011, January). Experimental investigation of the effect of rotor geometry and blade kinematics on the performance of a MAV-scale cycloidal rotor. In *Proceedings of the American Helicopter Society Specialists' Meeting on Unmanned Rotorcraft and Network Centric Operations*, Tempe, AZ.

[216] Benedict, M., Mattaboni, M., Chopra, I., and Masarati, P. (2011). Aeroelastic analysis of a micro-air-vehicle-scale cycloidal rotor in hover. *AIAA journal*, *49*(11), 2430-2443.

[217] Benedict, M., Ramasamy, M., and Chopra, I. (2010). Improving the aerodynamic performance of micro-air-vehicle-scale cycloidal rotor: an experimental approach. *Journal of aircraft*, *47*(4), 1117-1125.

[218] Benedict, M., Ramasamy, M., Chopra, I., and Leishman, J. G. (2009, January). Experiments on the optimization of the MAV-scale cycloidal rotor characteristics towards improving their aerodynamic performance. In *Proceedings of the International Specialists Meeting on Unmanned Rotorcraft*, Scottsdale, AZ (pp. 20-22).

[219] Benedict, M., Ramasamy, M., Chopra, I., and Leishman, J. G. (2010). Performance of a cycloidal rotor concept for micro air vehicle applications. *Journal of the American Helicopter Society*, *55*(2), 22002-22002.

[220] Bhattacharya, S., and Kumar, V. (2010). Distributed optimization with pairwise constraints and its application to multi-robot path planning.

[221] Bhattacharya, S., Likhachev, M., and Kumar, V. (2010, May). Multi-agent path planning with multiple tasks and distance constraints. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on* (pp. 953-959). IEEE.

[222] Bhattacharya, S., Michael, N., and Kumar, V. (2013). Distributed coverage and exploration in unknown non-convex environments. In *Distributed autonomous robotic systems* (pp. 61-75). Springer Berlin Heidelberg.

[223] Brockers, R., Bouffard, P., Ma, J., Matthies, L., and Tomlin, C. (2011, May). Autonomous landing and ingress of micro-air-vehicles in urban environments based on monocular vision. In *SPIE Defense, Security, and Sensing* (pp. 803111-803111). International Society for Optics and Photonics.

[224] Burden, S. A., Revzen, S., and Sastry, S. S. (2015). Model reduction near periodic orbits of hybrid dynamical systems. *IEEE Transactions on Automatic Control*, *60*(10), 2626-2639.

[225] Burden, S., Gonzalez, H., Vasudevan, R., Bajcsy, R., and Sastry, S. S. (2011, December). Numerical integration of hybrid dynamical systems via domain relaxation. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on* (pp. 3958-3965). IEEE.

[226] Bush, B., MacFarlane, K., Baeder, J., and Humbert, J. S. (2010, November). Development of immersed boundary code with application to mav stability analysis. In *27th Army Science Conference*, Orlando, FL.

[227] Cavagna, L., Masarati, P., and Quaranta, G. Coupled multibody/CFD Simulation of Maneuvering Flexible Aircraft. In *Journal of Aircraft*.

[228] Chang-Siu, E. H., Libby, T., Full, R. J., and Tomizuka, M. (2012, April). Tailbot-robot with inertial assisted control by an active tail inspired by lizards. *Integrative and Comparative Biology*, *52*, E224-E224. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[229] Chen, G., Sylvester, D., Blaauw, D., and Mudge, T. (2010). Yield-driven near-threshold SRAM design. *IEEE transactions on very large scale integration (VLSI) systems*, *18*(11), 1590-1598.

[230] Chen, G., Wieckowski, M., Blaauw, D., and Sylvester, D. (2010, September). Crosshairs SRAM—An adaptive memory for mitigating parametric failures. In *ESSCIRC, 2010 Proceedings of the* (pp. 366-369). IEEE.

[231] Chen, H., Hovareshti, P., and Baras, J. S. (2010, October). Opportunistic communications for networked controlled systems of autonomous vehicles. In *Military Communications Conference, MILCOM 2010* (pp. 1430-1435). IEEE.

[232] Chen, H., Hovareshti, P., and Baras, J. S. (2011, December). Distributed medium access and opportunistic scheduling for ad-hoc networks: an analysis of the constant access time problem. In *Global Telecommunications Conference (GLOBECOM 2011), 2011 IEEE* (pp. 1-6). IEEE.

[233] Choi, C., and Christensen, H. I. (2011, May). Robust 3D visual tracking using particle filtering on the SE (3) group. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on* (pp. 4384-4390). IEEE.

[234] Dagefu, F. T., Oh, J., and Sarabandi, K. (2013). A sub-wavelength RF source tracking system for GPS-denied environments. *IEEE Transactions on Antennas and Propagation*, *61*(4), 2252-2262.

[235] Dagefu, F., Oh, J., and Sarabandi, K. (2011, July). Numerical and experimental investigation of source tracking in a multipath environment. In *IEEE International Antennas and Propagation Symposium 2011*.

[236] DeMarco, K., West, M. E., and Collins, T. R. (2011, September). An implementation of ROS on the Yellowfin autonomous underwater vehicle (AUV). In *OCEANS 2011* (pp. 1-7). IEEE.

[237] Demir, A., Samson, E. G., Mongeau, J. M., Jayaram, K., Full, R. J., and Cowan, N. J. (2011, March). A tunable, multisegmented robotic antenna for identifying and testing biomechanical design principles. *Integrative and Comparative Biology*, *51*, E182-E182. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[238] Dickson, W. B, Polidoro, P. J., Tanner, M. M., and Dickinson, M. H. (2010, September). A linear systems analysis of the yaw dynamics of the fruit fly *Drosophila*. *J. Exp. Biol.*, 213(Pt 17), 3047-61.

[239] Ding, Y., Gravish, N., Li, C., Maladen, R. D., Mazouchova, N., Sharpe, S. S., and Goldman, D. I. (2012). Comparative studies reveal principles of movement on and within granular media. In *Natural Locomotion in Fluids and on Surfaces* (pp. 281-292). Springer New York.

[240] Diskin, B., and Yamaleev, N. K. (2011). Grid adaptation using adjoint-based error minimization. In *Proceedings of the 20th AIAA Computational Fluid Dynamics Conference* (No. 2011-3986).

[241] Elzinga, M. J., Dickson, W. B., and Dickinson, M. H. (2012). The influence of sensory delay on the yaw dynamics of a flapping insect. *Journal of The Royal Society Interface*, *9*(72), 1685-1696.

[242] Faruque, I., and Humbert, J. S. (2011, August). Reduced order roll/yaw model for dipteran flapping forward flight. In *AIAA Guidance, Navigation, and Control Conference* (p. 6548).

[243] Finio, B. M., and Wood, R. J. (2011, September). Optimal energy density piezoelectric twisting actuators. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 384-389). IEEE.

[244] Finio, B. M., Galloway, K. C., and Wood, R. J. (2011, September). An ultra-high precision, high bandwidth torque sensor for microrobotics applications. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 31-38). IEEE.

[245] Finio, B. M., Pérez-Arancibia, N. O., and Wood, R. J. (2011, September). System identification, modeling, and optimization of an insect-sized flapping-wing micro air vehicle. In *submitted to IEEE/RSJ Int. Conf. Intell. Robots Syst., San Francisco, CA*.

[246] Fink, J., Ribeiro, A., and Kumar, V. (2012). Robust control for mobility and wireless communication in cyber–physical systems with application to robot teams. *Proceedings of the IEEE*, *100*(1), 164-178.

[247] Fisher, T., Carpenter, M. H., Yamaleev, N. K., and Frankel, S. (2011). Boundary closures for 4th-order energy stable WENO finite difference schemes. *Journal of Computational Physics*, *230*, 3727-3752.

[248] Flynn, M. P., Lin, D. T., Ghahramani, M., and Li, L. (2011, May). New techniques for efficient flexible wireless transceivers in nanometer CMOS. In *SPIE Defense, Security, and Sensing* (pp. 803114-803114). International Society for Optics and Photonics.

[249] Forrest, S. R., Lee, K., and Zimmerman, J. (2010). Multiple growths of III-V solar cells on a single substrate using epitaxial lift-off. In *MRS Workshop on PV Materials and Manufacturing Issues*, Paper B5.1 (invited), Denver, CO. Oct. 6, 2010.

[250] Full, R. J., Jayaram, K., Mongeau, J. M., Birkmeyer, P., Hoover, A., and Fearing, R. S. (2011, March). Role of robustness in running: Bio-and bio-inspired exoskeletons. *Integrative and Comparative Biology*, *51*, E44-E44. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[251] Gerardi, S. A., Humbert, J. S., Pierce, L. E., and Sarabandi, K. (2011, May). Velocity estimation using optic flow and radar. In *SPIE Defense,*

Security, and Sensing (pp. 803112-803112). International Society for Optics and Photonics.

[252] Ghaffarkhah, A., and Mostofi, Y. (2010, June). Channel learning and communication-aware motion planning in mobile networks. In *American Control Conference (ACC), 2010* (pp. 5413-5420). IEEE.

[253] Ghaffarkhah, A., and Mostofi, Y. (2011). Communication-aware motion planning in mobile networks. *IEEE Transactions on Automatic Control*, *56*(10), 2478-2485.

[254] Ghaffarkhah, A., and Mostofi, Y. (2011, December). A communication-aware framework for robotic field estimation. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on* (pp. 3553-3558). IEEE.

[255] Ghaffarkhah, A., and Mostofi, Y. (2011, June). Communication-aware surveillance in mobile sensor networks. In *American Control Conference (ACC), 2011* (pp. 4032-4038). IEEE.

[256] Gil, S., Prentice, S., Roy, N., and Rus, D. (2017). Decentralized control for optimizing communication with infeasible regions. In *Robotics Research* (pp. 363-381). Springer International Publishing.

[257] Gokhale, V. J., Roberts, J., and Rais-Zadeh, M. (2011, June). High performance bulk mode gallium nitride resonators and filters. In *2011 16th International Solid-State Sensors, Actuators and Microsystems Conference* (pp. 926-929). IEEE.

[258] Hernández, E., Carreras, M., Galceran, E., and Ridao, P. (2011, June). Path planning with homotopy class constraints on bathymetric maps. In *OCEANS, 2011 IEEE-Spain* (pp. 1-6). IEEE.

[259] Hovareshti, P., Chen, H., and Baras, J. S. (2011). Motif-based topology design for efficient performance by networks of mobile autonomous vehicles.

[260] Hrishikeshavan, V., and Chopra, I. (2011). Aeromechanics and control of a shrouded rotor micro air vehicle in hover and in edgewise flow. *Journal of the American Helicopter Society*, *56*(4), 1-14.

[261] Hrishikeshavan, V., and Chopra, I. (2011, January). performance and rotor-based control of a shrouded rotor MAV in edgewise flow, In *AHS International Specialists Meeting on Unmanned Rotorcraft*, Tempe, AZ.

[262] Hrishikeshavan, V., and Chopra, I. (2011, July). Performance, control authority and flight testing of a shrouded rotor MAV in edgewise gust. *Journal of Aircraft*.

[263] Hrishikeshavan, V., Sirohi, J., Tishchenko, M., and Chopra, I. (2011). Design, development, and testing of a shrouded single-rotor micro air

vehicle with antitorque vanes. *Journal of the American Helicopter Society*, *56*(1), 12008-12008.

[264] Huang, A. S., Bachrach, A., Henry, P., Krainin, M., Maturana, D., Fox, D., and Roy, N. (2017). Visual odometry and mapping for autonomous flight using an RGB-D camera. In *Robotics Research* (pp. 235-252). Springer International Publishing.

[265] Fink, J., Ribeiro, A., Kumar, V., and Brian, M. S. (2010, October). Optimal robust multihop routing for wireless networks of mobile micro autonomous systems. In *Military Communications Conference, MILCOM 2010* (pp. 1268-1273). IEEE.

[266] Jayaram, K., Merritt, C., and Full, R. J. (2012, April). Robust climbing in cockroaches results from fault tolerant design using leg spines. *Integrative and Comparative Biology*, *52*, E86-E86. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[267] Jayaram, K., Merritt, C., Cherian, A., and Full, R. J. (2011, March). Running without feet: the role of tarsi during high-speed horizontal locomotion in cockroaches. *Integrative and Comparative Biology*, *51*, E64-E64. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[268] Jeon, D., Seok, M., Chakrabarti, C., Blaauw, D., and Sylvester, D. (2011, May). Energy-optimized high performance FFT processor. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1701-1704). IEEE.

[269] Jie, T., and Abbeel, P. (2010). On a connection between importance sampling and the likelihood ratio policy gradient. In *Advances in Neural Information Processing Systems* (pp. 1000-1008).

[270] Jones, M., and Yamaleev, N. K. (2012). The effect of a gust on the flapping wing performance. *AIAA*, *1080*, 2012.

[271] Jung, H., Menon, A., and Arkin, R. C. (2011). A conceptual space architecture for widely heterogeneous robotic systems.

[272] Jusufi, A., Kawano, D. T., Libby, T., and Full, R. J. (2010). Righting and turning in mid-air using appendage inertia: reptile tails, analytical models and bio-inspired robots. *Bioinspiration & Biomimetics*, *5*(4), 045001.

[273] Jusufi, A., Libby, T., and Full, R. J. (2012, April). Scales assist scaling in lizards: keeled, subcaudal scale arrays engage substrate during rapid vertical climbing. *Integrative and Comparative Biology*, *52*, E88-E88. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[274] Jusufi, A., Zeng, Y., Full, R. J., and Dudley, R. (2011). Aerial righting reflexes in flightless animals. *Integrative and Comparative Biology*, *51*(6), 937-943.

[275] Jusufi, A., Zeng, Y., Full, R. J., and Dudley, R. (2011, March). Aerial righting responses: a comparative approach. *Integrative and Comparative Biology*, *51*, E66-E66. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[276] Kallem, V., Komoroski, A. T., and Kumar, V. (2011). Sequential composition for navigating a nonholonomic cart in the presence of obstacles. *IEEE Transactions on Robotics*, *27*(6), 1152-1159.

[277] Kim, D., Chandra, V., Aitken, R., Blaauw, D., and Sylvester, D. (2011, August). Variation-aware static and dynamic writability analysis for voltage-scaled bit-interleaved 8-T SRAMs. In *Proceedings of the 17th IEEE/ACM international symposium on Low-power electronics and design* (pp. 145-150). IEEE Press.

[278] Lakshminarayan, V. K., and Baeder, J. D. (2011). Computational investigation of microscale shrouded rotor aerodynamics in hover. *Journal of the American Helicopter Society*, *56*(4), 1-15.

[279] Lee, K., Shiu, K. T., Zimmerman, J. D., Renshaw, C. K., and Forrest, S. R. (2010). Multiple growths of epitaxial lift-off solar cells from a single InP substrate. *Applied Physics Letters*, *97*(10), 101107.

[280] Lee, K., Zimmerman, J., and Forrest, S. R. (2011). Multiple growths of III-V solar cells from a single wafer. In *Global Photovoltaic Conference*, Busan, Korea.

[281] Lee, V., Sis, S. A., and Mortazawi, A. (2012). Ferroelectric resonators and filters for reconfigurable radios. In *Army Science Conference*.

[282] Libby, T., and Full, R. J. (2011, March). Active muscle enhances rapid perturbation recovery in an insect limb. *Integrative and Comparative Biology*, *51*, E81-E81. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[283] Libby, T., Moore, T. Y., Chang-Siu, E., Li, D., Cohen, D. J., Jusufi, A., and Full, R. J. (2012). Tail-assisted pitch control in lizards, robots and dinosaurs. *Nature*, *481*(7380), 181-184.

[284] Libby, T., Moore, T. Y., Chang-Siu, E., Li, D., Jusufi, J., Cohen, D., and Full, R. J. (2012, April). Tail assisted pitch control in a lizard, robot, and dinosaur. *Integrative and Comparative Biology*, *52*, E106-E106. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[285] Lin, D. T., Li, L., Farahani, S., and Flynn, M. P. (2012). A flexible 500 MHz to 3.6 GHz wireless receiver with configurable DT FIR and IIR filter embedded in a 7b 21 MS/s SAR ADC. *IEEE Transactions on Circuits and Systems I: Regular Papers*, *59*(12), 2846-2857.

[286] Lingaiah, S., and Shivakumar, K. (2013). Electrospun high temperature polyimide nanopaper. *European Polymer Journal*, *49*(8), 2101-2108.

[287] Lingaiah, S., Shivakumar, K., and Sadler, R. (2014). Electrospun nanopaper and its applications to microsystems. *International Journal for Computational Methods in Engineering Science and Mechanics*, *15*(1), 2-8.

[288] Macfarlane M., Bush B., Humbert J. S., and Baeder J. (2011, June). Estimation of a longitudinal model for hovering *Drosophila* using immersed boundary CFD. In *29th Applied Aerodynamics Conference*, Honolulu, HI.

[289] MacFarlane, K., Bush, B., Faruque, I., Humbert, J. S., and Baeder, J. (2011, June). Quasi-steady and computational aerodynamics applied to hovering Drosophila dynamics. In *AIAA 29th Applied Aerodynamics Conference*, Honolulu, HI.

[290] Maladen, R. D., Ding, Y., Umbanhowar, P. B., and Goldman, D. I. (2011). Undulatory swimming in sand: experimental and simulation studies of a robotic sandfish. *The International Journal of Robotics Research*, *30*(7), 793-805.

[291] Maladen, R. D., Ding, Y., Umbanhowar, P. B., Kamor, A., and Goldman, D. I. (2011). Mechanical models of sandfish locomotion reveal principles of high performance subsurface sand-swimming. *Journal of The Royal Society Interface*, *8*(62), 1332-1345.

[292] Maladen, R. D., Umbanhowar, P. B., Ding, Y., Masse, A., and Goldman, D. I. (2011, May). Granular lift forces predict vertical motion of a sand-swimming robot. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on* (pp. 1398-1403). IEEE.

[293] Malhan, R., Benedict, M., and Chopra, I. (2010, May). Experimental investigation of an avian-based flapping wing concept in hover and forward flight for micro air vehicle applications. In *66th Annual Forum of the American Helicopter Society*, Phoenix, AZ.

[294] Malhan, R., Benedict, M., and Chopra, I. (2012). Experimental studies to understand the hover and forward flight performance of a MAV-scale flapping wing concept. *Journal of the American Helicopter Society*, *57*(2), 1-11.

[295] Malhan, R., Lakshminarayan, V. K., Baeder, J., and Chopra, I. (2011, January). Investigation of aerodynamics of rigid flapping wings for MAV applications: CFD validation. In *AHS Specialists Conference* (pp. 25-27).

[296] Masarati, P., and Sitaraman, J. (2011, January). Coupled CFD/multibody analysis of NREL unsteady aerodynamic experiment phase VI rotor. In *AIAA Aerospace Sciences Meeting and Exhibit*, Orlando, FL. AIAA Paper (Vol. 153).

[297] Masarati, P., Morandini, M., Quaranta, G., and Vescovini, R. (2011). Multibody analysis of a micro-aerial vehicle flapping wing. *Multibody Dynamics*, 4-7.

[298] Mattaboni, M., Benedict, M., Masarati, P., and Chopra, I. (2010). MAV-scale cycloidal rotor multibody aeroelastic analysis. In *IMSD 2010*, Lappeenranta, Finland, May 25-27, 2010.

[299] Mayo, D. B., Ramsey, J. P., Leishman, J. G. (2012). Lift-to-drag ratio assessment of a flapping-wing using particle image velocimetry. In *30th Applied Aerodynamics Conference*, New Orleans, LA, June 25–28, 2012.

[300] Mehta, M., and Pister, K. (2011, October). Microrobots: Mobility in Wireless Sensor Networks. In *the 8th annual Graduate Students Academic Forum*, Beijing, China.

[301] Mellinger, D., Michael, N., and Kumar, V. (2012). Trajectory generation and control for precise aggressive maneuvers with quadrotors. *The International Journal of Robotics Research*, *31*(5), 664-674.

[302] Mellinger, D., Shomin, M., Michael, N., and Kumar, V. (2013). Cooperative grasping and transport using multiple quadrotors. In *Distributed autonomous robotic systems* (pp. 545-558). Springer Berlin Heidelberg.

[303] Michael, N., Schwager, M., Kumar, V., and Rus, D. (2014). An experimental study of time scales and stability in networked multi-robot systems. In *Experimental Robotics* (pp. 631-643). Springer Berlin Heidelberg.

[304] Mongeau, J. M., Alexander, T., and Full, R. J. (2012, April). Neuromechanical feedback during dynamic recovery after a lateral perturbation in rapid running cockroaches. *Integrative and Comparative Biology*, *52*, E297-E297. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[305] Mongeau, J. M., Jayaram, K., Demir, A., Sampson, E. G., Cowan, N. J., and Rull, R. J. (2011, March). Biomechanics of tactile sensor for wall following and spatial mapping. *Integrative and Comparative Biology*, *51*, E93-E93. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[306] Mongeau, J. M., McRae, B., Jusufi, A., Birkmeyer, P., Hoover, A. M., Fearing, R., and Full, R. J. (2012, April). Rapid inversion: running cockroaches, geckos, and robots swing like a pendulum under ledges. In *Integrative and Comparative Biology*, *52*, E123-E123. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[307] Mongeau, J. M., Sponberg, S. N., Miller, J. P., and Full, R. J. (2010). Revealing the neural controller in antenna-mediated tactile navigation. In *International Neuroethology Congress*, Salamanca, Spain.

[308] Mostofi, Y., Malmirchegini, M., and Ghaffarkhah, A. (2010, May). Estimation of communication signal strength in robotic networks. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on* (pp. 1946-1951). IEEE.

[309] Naroditsky, O., Zhou, X. S., Gallier, J., Roumeliotis, S. I., and Daniilidis, K. (2012). Two efficient solutions for visual odometry using directional correspondence. *IEEE transactions on pattern analysis and machine intelligence*, *34*(4), 818-824.

[310] Oh, J., and Sarabandi, K. (2011, August). A low-profile omnidirectional planar antenna with vertical polarization employing two in-phase elements. In *General Assembly and Scientific Symposium, 2011 XXXth URSI* (pp. 1-4). IEEE.

[311] Oh, J., and Sarabandi, K. (2011, July). A novel approach for miniaturization of circularly polarized patch antennas. In *Antennas and Propagation (APSURSI), 2011 IEEE International Symposium on* (pp. 880-883). IEEE.

[312] Oh, J., and SarabandI, K. (2012). Low profile, miniaturized, inductively coupled capacitively loaded monopole antenna. *IEEE Transactions on Antennas and Propagation*, *60*(3), 1206-1213.

[313] Panagou, D., Tanner, H. G., and Kyriakopoulos, K. J. (2011, December). Control of nonholonomic systems using reference vector fields. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on* (pp. 2831-2836). IEEE.

[314] Park, Y., and Wentzloff, D. D. (2011). A cyclic Vernier TDC for ADPLLs synthesized from a standard cell library. *IEEE Transactions on Circuits and Systems I: Regular Papers*, *58*(7), 1511-1517.

[315] Park, Y., and Wentzloff, D. D. (2011, September). An all-digital PLL synthesized from a digital standard cell library in 65nm CMOS. In *2011 IEEE Custom Integrated Circuits Conference (CICC)* (pp. 1-4). IEEE.

[316] Pathak, K., and Yamaleev, N. K. (2011). POD-based reduced-order model for arbitrary Mach number flows. In *6th AIAA Theoretical Fluid Mechanics Conference* (p. 3111).

[317] Peterson, K., and Fearing, R. S. (2011, September). Experimental dynamics of wing assisted running for a bipedal ornithopter. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 5080-5086). IEEE.

[318] Peterson, K., Birkmeyer, P., Dudley, R., and Fearing, R. S. (2011). A wing-assisted running robot and implications for avian flight evolution. *Bioinspiration & Biomimetics*, *6*(4), 046008.

[319] Peterson, K., Birkmeyer, P., Dudley, R., and Fearing, R. S. (2011, March). Wing assisted locomotion of a 25 g running robot. *Integrative and Comparative Biology*, *51*, E108-E108. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[320] Pierce, L., Humbert, S., Gerardi, S., Sarbandi, K., and Moallem, M. (2011, July). Combined radar and optical flow navigation for flying robots in urban areas. In *IEEE URSI/AP Conference*, Spokane, WA.

[321] Rogers III, J. G., Trevor, A. J., Nieto-Granda, C., Cunningham, A., Paluri, M., Michael, N., and Kumar, V. (2014). Effects of sensory precision on mobile robot localization and mapping. In *Experimental Robotics* (pp. 433-446). Springer Berlin Heidelberg.

[322] Rogers III, J. G., Trevor, A. J., Nieto-Granda, C., Cunningham, A., Paluri, M., Michael, N., and Kumar, V. (2014). Effects of sensory precision on mobile robot localization and mapping. In *Experimental Robotics* (pp. 433-446). Springer Berlin Heidelberg.

[323] Rogers, J. G., Trevor, A. J., Nieto-Granda, C., and Christensen, H. I. (2011, September). Simultaneous localization and mapping with learned object recognition and semantic data association. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on* (pp. 1264-1270). IEEE.

[324] Rogers, J., Cunningham, A., Paluri, M., Christensen, H. I., Michaels, N., Kumar, V., and Matthies, L. (2011, April). Multi-robot mapping. In *SPIE Defense XIII*, Florida.

[325] Sadeghi, M. M., Peterson, R. L., and Najafi, K. (2011, December). Micro-hydraulic structure for high performance bio-mimetic air flow sensor arrays. In *Electron Devices Meeting (IEDM), 2011 IEEE International* (pp. 29-4). IEEE.

[326] Sadeghi, M., Peterson, R. L., Peterson, K., Fearing, R., and Najafi, K. (2010, December). Air-flow sensing on autonomous mobile platforms using micro-scale hot-wire anemometry. In *27th Army Science Conference*.

[327] Sarabandi, K., and Song, Y. J. (2011). Subwavelength radio repeater system utilizing miniaturized antennas and metamaterial channel isolator. *IEEE Transactions on Antennas and Propagation*, *59*(7), 2683-2690.

[328] Sarabandi, K., Song, Y. J., and Oh, J. (2011, May). Enhanced ad hoc wireless connectivity in complex environment using small radio repeater systems. In *SPIE Defense, Security, and Sensing* (pp. 803117-803117). International Society for Optics and Photonics.

[329] Sciarretta, A., Mait, J., Chait, R., Redden, E., and Wilcox, J. (2011). *Assessing military benefits of S&T investments in micro autonomous*

systems utilizing a gedanken experiment. National Defense Univ Washington DC Center for Technology and National Security Policy.

[330] Seok, M., Jeon, D., Chakrabarti, C., Blaauw, D., and Sylvester, D. (2011, February). A 0.27 V 30MHz 17.7 nJ/transform 1024-pt complex FFT core with super-pipelining. In *2011 IEEE International Solid-State Circuits Conference* (pp. 342-344). IEEE.

[331] Seok, M., Jeon, D., Chakrabarti, C., Blaauw, D., and Sylvester, D. (2011, June). Pipeline strategy for improving optimal energy efficiency in ultra-low voltage design. In *Design Automation Conference (DAC), 2011 48th ACM/EDAC/IEEE* (pp. 990-995). IEEE.

[332] Seshadri, P., Benedict, M., and Chopra, I. (2011). Aerodynamics and control towards a biomimetic hovering MAV. In *AIAA Region I-MA Student Conference*, Charlottesville, VA, April 8-9, 2011.

[333] Seshadri, P., Benedict, M., and Chopra, I. (2011). Control of a biomimetic flapping mechanism for a hovering micro air vehicle. In *Proceedings of the American Helicopter Society International Specialists' Meeting on Unmanned Systems and Network Centric Operations*, Tempe, AZ, January 25-27, 2011.

[334] Shah, S. K., and Tanner, H. G. (2011, June). Bounding the uncertainity in nonlinear robust model predictive control using sphere covering. In *Control & Automation (MED), 2011 19th Mediterranean Conference on* (pp. 807-812). IEEE.

[335] Shah, S. K., Pahlajani, C. D., and Tanner, H. G. (2011, September). Probability of success in stochastic robot navigation with state feedback. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on* (pp. 3911-3916). IEEE.

[336] Shen, S., Michael, N., and Kumar, V. (2011, May). Autonomous multi-floor indoor navigation with a computationally constrained MAV. In *Robotics and automation (ICRA), 2011 IEEE international conference on* (pp. 20-25). IEEE.

[337] Shenderova, O. A., Hens, S., Sumant, A. V., Yu, S., Gokhale, V. J., Rais-Zadeh, M., and McGuire, G. E. (2011, September). Detonation nanodiamond in DMSO as seeding slurries for CVD diamond. In *IV International Conference on Surfaces, Materials and Vacuum*, Puerto Vallarta, Mexico.

[338] Shivakumar, K. N., Chen, H., and Ali, M. (2012). Nanofiber interleaved composites for aerospace applications. *Journal of Aerospace Science and Technologies*, *64*(1), 55.

[339] Shivakumar, K., and Swaminathan, G. (2011). Thermoset polymer nanocomposites: processing, fabrication and characterization, a short course. In *ICC-CFT-2011 Conference*, Bangalore, India, January 4-7, 2011.

[340] Sis, S. A., Lee, V., and Mortazawi, A. (2011, June). Intrinsically switchable, BST-on-silicon composite FBARs. In *Microwave Symposium Digest (MTT), 2011 IEEE MTT-S International* (pp. 1-4). IEEE.

[341] Straw, A. D., Lee, S., and Dickinson, M. H. (2010). Visual control of altitude in flying *Drosophila*. *Current Biology*, *20*(17), 1550-1556.

[342] Stump, E., Michael, N., Kumar, V., and Isler, V. (2011, May). Visibility-based deployment of robot formations for communication maintenance. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on* (pp. 4498-4505). IEEE.

[343] Talucdher, A. R., Sergey, Y., Lingaiah, S., and Shivakumar, K. (2011). Indentation properties of NC dragonfly wing by nanolndentation. In *26th Annual Technical Conference of the American Society for Composites 2011* and the *2nd Joint US-Canada Conference on Composites*.

[344] Whitney, J. P., Sreetharan, P. S., Ma, K. Y., and Wood, R. J. (2011). Pop-up book MEMS. *Journal of Micromechanics and Microengineering*, *21*(11), 115021.

[345] Williams, B., Hudson, N., Tweddle, B., Brockers, R., and Matthies, L. (2011, May). Feature and pose constrained visual aided inertial navigation for computationally constrained aerial vehicles. In *Robotics and automation (ICRA), 2011 IEEE international conference on* (pp. 431-438). IEEE.

[346] Zavlanos, M. M., Ribeiro, A., and Pappas, G. J. (2010, December). Mobility & routing control in networks of robots. In *Decision and Control (CDC), 2010 49th IEEE Conference on* (pp. 7545-7550). IEEE.

FY12

[347] Aguilar, J., Lesov, A., Wiesenfeld, K., and Goldman, D. I. (2012). Lift-off dynamics in a simple jumping robot. *Physical review letters*, *109*(17), 174301.

[348] Arkin, R. C., Garcia-Vergara, S., and Lee, S. G. (2012, May). Architectural design and support for knowledge sharing across heterogeneous MAST systems. In *SPIE Defense, Security, and Sensing* (pp. 84070C-84070C). International Society for Optics and Photonics.

[349] Atanasov, N., Le Ny, J., Michael, N., and Pappas, G. J. (2012, May). Stochastic source seeking in complex environments. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on* (pp. 3013-3018). IEEE.

[350] Bachrach, A., Prentice, S., He, R., Henry, P., Huang, A. S., Krainin, M., and Roy, N. (2012). Estimation, planning, and mapping for autonomous flight using an RGB-D camera in GPS-denied environments. *The International Journal of Robotics Research*, *31*(11), 1320-1343.

[351] Benedict, M., and Chopra, I. (2012, May). Design and development of an unconventional VTOL micro air vehicle: The Cyclocopter. In *SPIE Defense, Security, and Sensing* (pp. 83731F-83731F). International Society for Optics and Photonics.

[352] Benedict, M., Gupta, R., and Chopra, I. (2011). Design, development and flight testing of a twin-rotor cyclocopter micro air vehicle. In *Proceedings of the 67th Annual National Forum of the American Helicopter Society*, Virginia Beach, VA, May 3–5, 2011.

[353] Benedict, M., Gupta, R., and Chopra, I. (2013). Design, development, and open-loop flight-testing of a twin-rotor cyclocopter micro air vehicle. *Journal of the American Helicopter Society*, *58*(4), 1-10.

[354] Benedict, M., Jarugumilli, T., and Chopra, I. (2011, January). Experimental investigation of the effect of rotor geometry and blade kinematics on the performance of a MAV-scale cycloidal rotor. In *Proceedings of the American Helicopter Society Specialists' Meeting on Unmanned Rotorcraft and Network Centric Operation*s, Tempe, AZ.

[355] Benedict, M., Jarugumilli, T., Lakshminarayan, V., and Chopra, I. (2014). Effect of flow curvature on forward flight performance of a micro-air-vehicle-scale cycloidal-rotor. *AIAA journal*, *52*(6), 1159-1169.

[356] Benedict, M., Jarugumilli, T., Lakshminarayan, V., and Chopra, I. (2012, April). Experimental and computational studies to understand the role of flow curvature effects on the aerodynamic performance of a MAV-scale cycloidal rotor in forward flight. In *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference 20th AIAA/ASME/AHS Adaptive Structures Conference 14th AIAA* (p. 1629).

[357] Benedict, M., Mattaboni, M., Chopra, I., and Masarati, P. (2011). Aeroelastic analysis of a micro-air-vehicle-scale cycloidal rotor in hover. *AIAA journal*, *49*(11), 2430-2443.

[358] Benedict, M., Mullins, J., Hrishikeshavan, V., and Chopra, I. (2013). Development of an 800 gram quad-cyclocopter capable of autonomous hover. Abstract at *the American Helicopter Society International Specialists' Meeting on Unmanned Rotorcraft*, Scottsdale, AZ, January 22-24, 2013.

[359] Benedict, M., Mullins, J., Hrishikeshavan, V., and Chopra, I. (2016). Development of a quad cycloidal-rotor unmanned aerial vehicle. *Journal of the American Helicopter Society*, *61*(2), 1-12.

[360] Benedict, M., Ramasamy, M., and Chopra, I. (2010). Improving the aerodynamic performance of micro-air-vehicle-scale cycloidal rotor: an experimental approach. *Journal of aircraft*, *47*(4), 1117-1125.

[361] Benedict, M., Ramasamy, M., Chopra, I., and Leishman, J. G. (2010). Performance of a cycloidal rotor concept for micro air vehicle applications. *Journal of the American Helicopter Society*, *55*(2), 22002-22002.

[362] Benedict, M., Shrestha, E., Hrishikeshavan, V., and Chopra, I. (2012, January). Development of 200 gram twin-rotor micro cyclocopter capable of autonomous hover. In *Proceedings of the American Helicopter Society Future Vertical Lift Aircraft Design Conference, San Francisco, CA* (pp. 18-20).

[363] Birkmeyer, P., Peterson, K., and Fearing, R. S. (2009, October). DASH: A dynamic 16g hexapedal robot. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on* (pp. 2683-2689). IEEE.

[364] Brockers, R., Susca, S., Zhu, D., and Matthies, L. (2012, May). Fully self-contained vision-aided navigation and landing of a micro air vehicle independent from external sensor inputs. In *SPIE Defense, Security, and Sensing* (pp. 83870Q-83870Q). International Society for Optics and Photonics.

[365] Bry, A., Bachrach, A., and Roy, N. (2012, May). State estimation for aggressive flight in GPS-denied environments using onboard sensing. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on* (pp. 1-8). IEEE.

[366] Burden, S., Ohlsson, H., and Sastry, S. S. (2012). Parameter identification near periodic orbits of hybrid dynamical systems. *IFAC Proceedings Volumes*, *45*(16), 1197-1202.

[367] Chan, W. R., Rinnerbauer, V., Yeng, Y. X., Senkevich, J., Soljacic, M., Joannopoulos, J. D., and Celanovic, I. (2012, September). Millimeter scale metal based catalytic reactor TPV system, In *27th EU PVSEC*, Frankfurt.

[368] Chang-Siu, E., Libby, T., Tomizuka, M., and Full, R. J. (2011, September). A lizard-inspired active tail enables rapid maneuvers and dynamic stabilization in a terrestrial robot. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on* (pp. 1887-1894). IEEE.

[369] Chang-Siu, E., T. Libby, M. Tomizuka, and R. J. Full. (2011, September). Lizards' tails inspire inertia assisted robotics for improved terrestrial locomotion. In *IEEE/RSJ International Conference on Intelligent Robots and Systems,* San Francisco, CA.

[370] Choi, J., Park, S., Cho, J., and Yoon, E. (2012, February). A 1.36μW adaptive CMOS image sensor with reconfigurable modes of operation from

available energy/illumination for distributed wireless sensor network. In *2012 IEEE International Solid-State Circuits Conference* (pp. 112-114). IEEE.

[371] Cunningham, A., and Dellaert, F. (2012, May). Large-scale experimental design for decentralized SLAM. In *SPIE Defense, Security, and Sensing* (pp. 83870O-83870O). International Society for Optics and Photonics.

[372] Cunningham, A., Indelman, V., and Dellaert, F. (2012, November). Consistent decentralized graphical SLAM with anti-factor down-dating. In *Safety, Security, and Rescue Robotics (SSRR), 2012 IEEE International Symposium on* (pp. 1-2). IEEE.

[373] Cunningham, A., Wurm, K. M., Burgard, W., and Dellaert, F. (2012, May). Fully distributed scalable smoothing and mapping with robust multi-robot data association. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on* (pp. 1093-1100). IEEE.

[374] Dagefu, F. T., Oh, J., and Sarabandi, K. (2013). A sub-wavelength RF source tracking system for GPS-denied environments. *IEEE Transactions on Antennas and Propagation*, *61*(4), 2252-2262.

[375] Dallmann, C. J., Mongeau, J. M., Jayaram, K., Mahavadi, A., and Full, R. J. (2013, April). Dynamic response of antenna flagellum in the American cockroach. In *Integrative and Comparative Biology*, *53*, E273-E273. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[376] Ding, Y., Gravish, N., and Goldman, D. I. (2011). Drag induced lift in granular media. *Physical Review Letters*, *106*(2), 028001.

[377] Ding, Y., Gravish, N., Li, C., Maladen, R. D., Mazouchova, N., Sharpe, S. S., and Goldman, D. I. (2012). Comparative studies reveal principles of movement on and within granular media. In *Natural Locomotion in Fluids and on Surfaces* (pp. 281-292). Springer New York.

[378] Ding, Y., Sharpe, S. S., Masse, A., and Goldman, D. I. (2012). Mechanics of undulatory swimming in a frictional fluid. *PLoS Comput Biol*, *8*(12), e1002810.

[379] Elzinga, M. J., Dickson, W. B., and Dickinson, M. H. (2012). The influence of sensory delay on the yaw dynamics of a flapping insect. *Journal of The Royal Society Interface*, *9*(72), 1685-1696.

[380] Faisal, M., Park, Y., and Wentzloff, D. D. (2011, May). Reconfigurable firmware-defined radios synthesized from standard digital logic cells. In *SPIE Defense, Security, and Sensing* (pp. 803115-803115). International Society for Optics and Photonics.

[381] Fick, D., Dreslinski, R. G., Giridhar, B., Kim, G., Seo, S., Fojtik, M., and Wieckowski, M. (2012, February). Centip3De: A 3930DMIPS/W

configurable near-threshold 3D stacked system with 64 ARM Cortex-M3 cores. In *2012 IEEE International Solid-State Circuits Conference* (pp. 190-192). IEEE.

[382] Fojtik, M., Fick, D., Kim, Y., Pinckney, N., Harris, D., Blaauw, D., and Sylvester, D. (2012, February). Bubble Razor: An architecture-independent approach to timing-error detection and correction. In *2012 IEEE International Solid-State Circuits Conference* (pp. 488-490). IEEE.

[383] Forrest, S., Lee, K., and Zimmerman, J.D. (2012). Epitaxial lift-off of GaAs thin film solar cells followed by substrate reuse. In *Materials Research Society 2012 Fall Meeting*, Boston, MA. November 26-30, 2012 (invited).

[384] Forrest, S.R. (2012). Electronics on 3D substrates. In *Materials Research Society 2012 Spring Meeting*, invited, Paper K6.6, San Francisco, CA, April 10-13, 2012.

[385] Ghaffarkhah, A., and Mostofi, Y. (2011, December). A communication-aware framework for robotic field estimation. In *2011 50th IEEE Conference on Decision and Control and European Control Conference* (pp. 3553-3558). IEEE.

[386] Ghaffarkhah, A., and Mostofi, Y. (2011, June). Communication-aware surveillance in mobile sensor networks. In *Proceedings of the 2011 American Control Conference* (pp. 4032-4038). IEEE.

[387] Ghaffarkhah, A., and Mostofi, Y. (2012). Path planning for networked robotic surveillance. *IEEE Transactions on Signal Processing*, *60*(7), 3560-3575.

[388] Ghaffarkhah, A., Yan, Y., and Mostofi, Y. (2011, December). Dynamic coverage of time-varying environments using a mobile robot—A communication-aware perspective. In *2011 IEEE GLOBECOM Workshops (GC Wkshps)* (pp. 1297-1302). IEEE.

[389] Ghosh, D., and Baeder, J. D. (2012). Compact reconstruction schemes with weighted ENO limiting for hyperbolic conservation laws. *SIAM Journal on Scientific Computing*, *34*(3), A1678-A1706.

[390] Ghosh, D., Baeder, J.D. (2012). High order compact reconstruction scheme with weighted essentially non-oscillatory limiting. In *AMS-MAA Joint Mathematics Meetings*, Boston, MA, Jan 4-7, 2012.

[391] Gokhale, V. J., Sui, Y., and Rais-Zadeh, M. (2012, May). Novel uncooled detector based on gallium nitride micromechanical resonators. In *SPIE Defense, Security, and Sensing* (pp. 835319-835319). International Society for Optics and Photonics.

[392] Gokhale, V.J., and Rais-Zadeh, M. (2012). Resonant uncooled IR detectors using gallium nitride micromechanical resonators and silicon nitride absorbers. In *Hilton Head Workshop*.

[393] Gordenker, R. J., and Wise, K. D. (2012, May). A programmable palm-size gas analyzer for use in micro-autonomous systems. In *SPIE Defense, Security, and Sensing* (pp. 83731O-83731O). International Society for Optics and Photonics.

[394] Gravish, N., Franklin, S. V., Hu, D. L., and Goldman, D. I. (2012). Entangled granular media. *Physical review letters*, *108*(20), 208001.

[395] Gravish, N., Garcia, M., Mazouchova, N., Levy, L., Umbanhowar, P. B., Goodisman, M. A., and Goldman, D. I. (2012). Effects of worker size on the dynamics of fire ant tunnel construction. *Journal of The Royal Society Interface*, rsif20120423.

[396] Gudenus, V., Libby, T., Haldane, D. W., and Full, R. J. (2013, April). Designing bio-inspired tailed robots to turn using bendable backs. In *Integrative and Comparative Biology*, *53*, E291-E291. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[397] Hatton, R. L., Ding, Y., Choset, H., and Goldman, D. I. (2013). Geometric visualization of self-propulsion in a complex medium. *Physical review letters*, *110*(7), 078101.

[398] Hatton, R. L., Ding, Y., Maladen, R. D., Masse, A., Choset, H., and Goldman, D. I. (2012). On the geometric mechanics of sand-swimming. In *Proceedings of the National Academy of Sciences*, in 2nd review.

[399] Hrishikeshavan, V., and Chopra, I. (2012). Performance, flight testing of a shrouded rotor micro air vehicle in edgewise gusts. *Journal of Aircraft*, *49*(1), 193-205.

[400] Hrishikeshavan, V., and Chopra, I. (2011, July). Performance, control authority and flight testing of a shrouded rotor MAV in edgewise gust. *Journal of Aircraft*.

[401] Jarugumilli, T., Benedict, M., and Chopra, I. (2012, May). Experimental investigation of the forward flight performance of a MAV-scale cycloidal rotor. In *Proceedings of the 68th Annual National Forum of the American Helicopter Society* (pp. 112-130). Fort Worth, TX: Curran Associates, Inc.

[402] Jayaram, K., Springthorpe, D., Haldane, D., Mckinley, S., Dirocco, A., and Full, R. J. (2013, April). Running in confined spaces by the American cockroach. In *Integrative and Comparative Biology*, *53*, E102-E102. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[403] Jeon, D., Seok, M., Chakrabarti, C., Blaauw, D., and Sylvester, D. (2012). A super-pipelined energy efficient subthreshold 240 MS/s FFT core in 65 nm CMOS. *IEEE Journal of Solid-State Circuits*, *47*(1), 23-34.

[404] Jones, M. and Yamaleev, N. K. (2012). Adjoint-based optimization of the flapping wing performance. In *Proc. of the 7th Inter. Conf. on CFD*, Island of Hawaii, HI, July 9-13, 2012.

[405] Jones, M., and Yamaleev, N. (2012). The effect of a gust on the flapping wing performance. In *50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition* (p. 1080).

[406] Jones, M., and Yamaleev, N. Adjoint-based optimization of the flapping wing performance.

[407] Jones., M., and Yamaleev, N. (2012). Adjoint-based optimization of the flapping wing performance. In *7th Inter. Conf. on CFD*, Island of Hawaii, HI, July 9-13, 2012.

[408] Jusufi A, Byrnes GT, Chang E, Libby T., and Full RJ (2011, in prep) Landing on a wall: arboreal lizards and acrobatic robots reveal common mechanism of tail-assisted perching.

[409] Jusufi, A., Byrnes, G. T., and Full, R. J. (2013, April). Gliding Geckos Perch on a Tree Trunk Assisted by Active Tails. In *Integrative and Comparative Biology*, *53*, E105-E105. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[410] Karydis, K., Poulakakis, I., and Tanner, H. G. (2012, October). A switching kinematic model for an octapedal robot. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on* (pp. 507-512). IEEE.

[411] Kohut, N. J., Birkmeyer, P. M., Peterson, K. C., and Fearing, R. S. (2012, May). Maneuverability and mobility in palm-sized legged robots. In *SPIE Defense, Security, and Sensing* (pp. 83731I-83731I). International Society for Optics and Photonics.

[412] Kohut, N. J., Haldane, D. W., Zarrouk, D., and Fearing, R. S. (2012, July). Effect of inertial tail on yaw rate of 45 gram legged robot. In *Int. Conf. Climbing Walk. Robot. Support Technol. Mob. Mach* (pp. 157-164).

[413] Kolluru, S., and Jones, A. R. (2012). Leading edge vortex development on a pitching and rotating wing. In *the 50th AIAA Aerospace Sciences Meeting and Exhibit*, Nashville, TN, Jan. 9-12, 2012. Paper 2012-0050.

[414] Kolluru, S., and Jones, A. R. Leading edge vortex structure over multiple revolutions of a rotating wing. In preparation for *Journal of Aircraft*.

[415] Kumar, S., Shi, L., Ahmed, N., Gil, S., Katabi, D., and Rus, D. (2012). Carspeak: a content-centric network for autonomous driving. *ACM SIGCOMM Computer Communication Review*, *42*(4), 259-270.

[416] Lakshminarayan, V. K., and Baeder, J. D. (2011). Computational investigation of microscale shrouded rotor aerodynamics in hover. *Journal of the American Helicopter Society*, *56*(4), 1-15.

[417] Lee, K., Zimmerman, J. D., Xiao, X., Sun, K., and Forrest, S. R. (2012). Reuse of GaAs substrates for epitaxial lift-off by employing protection layers. *Journal of Applied Physics*, *111*(3), 033527.

[418] Lee, K., Zimmerman, J. D., Zhang, Y., and Forrest, S. R. (2012, June). Epitaxial lift-off of GaAs thin-film solar cells followed by substrate reuse. In *Photovoltaic Specialists Conference (PVSC), 2012 38th IEEE* (pp. 001698-001700). IEEE.

[419] Lee, V., Lee, S., Sis, S. A., and Mortazawi, A. (2012, June). Large signal performance of ferroelectric FBARs. In *Microwave Symposium Digest (MTT), 2012 IEEE MTT-S International* (pp. 1-3). IEEE.

[420] Li, C., Ding, Y., Gravish, N., Maladen, R. D., Masse, A., Umbanhowar, P. B., and Goldman, D. I. (2012). Toward a terramechanics for bio-inspired locomotion in granular environments. In *Earth and Space 2012: Engineering, Science, Construction, and Operations in Challenging Environments* (pp. 264-273).

[421] Li, C., Hsieh, S. T., and Goldman, D. I. (2012). Multi-functional foot use during running in the zebra-tailed lizard (Callisaurus draconoides). *Journal of Experimental Biology*, *215*(18), 3293-3308.

[422] Li, C., Zhang, T., and Goldman, D. I. (2013). A terradynamics of legged locomotion on granular media. *Science*, *339*(6126), 1408-1412.

[423] Li, L., Flynn, M. P., and Ferriss, M. A. (2012, November). A 5.8 GHz digital arbitrary phase-setting Type II PLL in 65nm CMOS with 2.25° Resolution. In *Solid State Circuits Conference (A-SSCC), 2012 IEEE Asian* (pp. 317-320). IEEE.

[424] Libby, T., Gudenus, V., Harrington, P., and Full, R. J. (2013, April). Coordination between inertial and impulsive mechanisms during rapid turns in lizards. In *Integrative and Comparative Biology*, *53*, E127-E127. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[425] Libby, T., Moore, T. Y., Chang-Siu, E., Li, D., Cohen, D. J., Jusufi, A., and Full, R. J. (2012). Tail-assisted pitch control in lizards, robots and dinosaurs. *Nature*, *481*(7380), 181-184.

[426] Lin, D. T., Chae, H., Li, L., and Flynn, M. P. (2012, June). A 600 MHz to 3.4 GHz flexible spectrum-sensing receiver with spectrum-adaptive reconfigurable DT filtering. In *2012 IEEE Radio Frequency Integrated Circuits Symposium* (pp. 269-272). IEEE.

[427] Lin, D. T., Li, L., Farahani, S., and Flynn, M. P. (2012). A flexible 500 MHz to 3.6 GHz wireless receiver with configurable DT FIR and IIR filter embedded in a 7b 21 MS/s SAR ADC. *IEEE Transactions on Circuits and Systems I: Regular Papers*, *59*(12), 2846-2857.

[428] Lin, D. T., Li, L., Farahani, S., and Flynn, M. P. (2012). A flexible 500 MHz to 3.6 GHz wireless receiver with configurable DT FIR and IIR filter embedded in a 7b 21 MS/s SAR ADC. *IEEE Transactions on Circuits and Systems I: Regular Papers*, *59*(12), 2846-2857.

[429] Lingaiah, S., and Shivakumar, K. (2013). Electrospun high temperature polyimide nanopaper. *European Polymer Journal*, *49*(8), 2101-2108.

[430] Maladen, R. D., Ding, Y., Umbanhowar, P. B., and Goldman, D. I. (2011). Undulatory swimming in sand: experimental and simulation studies of a robotic sandfish. *The International Journal of Robotics Research*, *30*(7), 793-805.

[431] Maladen, R. D., Ding, Y., Umbanhowar, P. B., Kamor, A., and Goldman, D. I. (2011). Mechanical models of sandfish locomotion reveal principles of high performance subsurface sand-swimming. *Journal of The Royal Society Interface*, *8*(62), 1332-1345.

[432] Maladen, R. D., Umbanhowar, P. B., Ding, Y., Masse, A., and Goldman, D. I. (2011, May). Granular lift forces predict vertical motion of a sand-swimming robot. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on* (pp. 1398-1403). IEEE.

[433] Malhan, R., Benedict, M., and Chopra, I. (2012). Experimental studies to understand the hover and forward flight performance of a MAV-scale flapping wing concept. *Journal of the American Helicopter Society*, *57*(2), 1-11.

[434] Malhan, R., Lakshminarayan, V. K., Baeder, J., and Chopra, I. (2011, January). Investigation of aerodynamics of rigid flapping wings for MAV applications: CFD validation. In *AHS Specialists Conference* (pp. 25-27).

[435] Malhan, R., Lakshminarayan, V., Baeder, J., Chopra, I., Masarati, P., Morandini, M., and Quaranta, G. (2012). CFD-CSD coupled aeroelastic analysis of flexible flapping wings for MAV applications: methodology validation. In *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference 20th AIAA/ASME/AHS Adaptive Structures Conference 14th AIAA* (p. 1636).

[436] Malhan, R., Lakshminarayan, V., Masarati, P., Baeder, J., and Chopra, I. (2012). Investigation of aerodynamics of flexible flapping wings for MAV applications: methodology validation. In *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Honolulu, Hawaii, 23-26 April 2012.

[437] Malmirchegini, M., and Mostofi, Y. (2012). An integrated sparsity and model-based probabilistic framework for estimating the spatial variations of communication channels. *Physical Communication*, *5*(2), 102-118.

[438] Malmirchegini, M., and Mostofi, Y. (2012). On the spatial predictability of communication channels. *IEEE Transactions on Wireless Communications*, *11*(3), 964-978.

[439] Masarati P., Quaranta, G., Roget, B., and Sitaraman, J. (2011). Unsteady CFD/CSD analysis of self-activated aerodynamic devices for stall alleviation. In *IFASD 2011*, Paris.

[440] Masarati, M., Morandini M., Quaranta, G., Chandar, D., Roget, B., and Sitaraman, J. (2011). Tightly coupled CFD/multibody analysis of flapping-wing MAV. In *41st AIAA Fluid Dynamics Conference*, 2011, Honolulu, HI.

[441] Masarati, P., Morandini, M., Quaranta, G., and Vescovini, R. (2011). Multibody analysis of a micro-aerial vehicle flapping wing. *Multibody Dynamics, JC Samin and P. Fisette, eds., Brussels, Belgium*.

[442] Masarati, P., Morandini, M., Quaranta, G., Chandar, D., Roget, B., and Sitaraman, J. (2011). Tightly coupled CFD/multibody analysis of flapping-wing micro-aerial vehicles. In *29th AIAA Applied Aerodynamics Conference* (p. 3022).

[443] Mayo, D. (2012). Flow visualization and PIV measurements of a leading-edge vortex generated by a rotating wing. In *2012 NOBCChE Conference* September 23-25, 2012.

[444] Mayo, D. B., and Leishman J. G. (2012). PIV measurements of the leading-edge vortex generated by a flapping-wing MAV. In *AHS 68th Annual Forum*, Fort Worth, TX, May 1–3, 2012.

[445] Mayo, D. B., Ramsey, J. P., and Leishman, J. G. (2012). Lift-to-drag ratio assessment of a flapping-wing using particle image velocimetry. In *30th Applied Aerodynamics Conference*, New Orleans, LA, June 25–28, 2012.

[446] Mayo, D., Kolluru, S., Glucksman-Glaser, M., and Jones, A. R. (2013). The effect of wake encounters on the leading edge vortex generated by a rotating wing. In *the 51st AIAA Aerospace Sciences Meeting and Exhibit*, Grapevine, TX, Jan. 7-10, 2013.

[447] Mazouchova, N., Umbanhowar, P. B., and Goldman, D. I. (2013). Flipper-driven terrestrial locomotion of a sea turtle-inspired robot. *Bioinspiration & biomimetics*, *8*(2), 026007.

[448] Michael, N., Shen, S., Mohta, K., Mulgaonkar, Y., Kumar, V., Nagatani, K., and Ohno, K. (2012). Collaborative mapping of an earthquake-damaged building via ground and aerial robots. *Journal of Field Robotics*, *29*(5), 832-841.

[449] Miller, L. A., Goldman, D. I., Hedrick, T. L., Tytell, E. D., Wang, Z. J., Yen, J., and Alben, S. (2012). Using computational and mechanical models to study animal locomotion. *Integrative and comparative biology*, ics115.

[450] Moldavan, T., and Abbeel, P. (2012). Near optimal Chernoff bounds for Markov decision processes. In *Neural Information Processing Systems (NIPS)*.

[451] Moldovan, T. M., and Abbeel, P. (2012, December). Risk aversion in Markov decision processes via near optimal Chernoff bounds. In *NIPS* (pp. 3140-3148).

[452] Mongeau, J. M., McRae, B., Jusufi, A., Birkmeyer, P., Hoover, A. M., Fearing, R., and Full, R. J. (2012). Rapid inversion: Running animals and robots swing like a pendulum under ledges. *PloS one*, *7*(6), e38003.

[453] Mongeau, J. M., Sponberg, S. N., and Full, R. J. (2013, April). Unit responses from antenna in cockroaches generate control input predicted from control-theoretic model of wall following. In *Integrative and Comparative Biology*, *53*, E148-E14). Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[454] Naroditsky, O., Zhou, X. S., Gallier, J., Roumeliotis, S. I., and Daniilidis, K. (2012). Two efficient solutions for visual odometry using directional correspondence. *IEEE transactions on pattern analysis and machine intelligence*, *34*(4), 818-824.

[455] Oh, J., and SarabandI, K. (2012). Low profile, miniaturized, inductively coupled capacitively loaded monopole antenna. *IEEE Transactions on Antennas and Propagation*, *60*(3), 1206-1213.

[456] Oh, J., and Sarabandi, K. (2012, July). Ultra-wideband, miniaturized, low profile, omnidirectional antenna using a novel reactive loading approach. In *Proceedings of the 2012 IEEE International Symposium on Antennas and Propagation* (pp. 1-2). IEEE.

[457] Oh, J., and Sarabandi, K. (2012, March). A moderate gain extremely short HF monopole antenna. In *2012 6th European Conference on Antennas and Propagation (EUCAP)* (pp. 1657-1660). IEEE.

[458] Oh, J., and Sarabandi, K. (2013). A topology-based miniaturization of circularly polarized patch antennas. *IEEE Transactions on Antennas and Propagation*, *61*(3), 1422-1426.

[459] Oh, J., and Sarabandi, K. (2014). Low profile vertically polarized omnidirectional wideband antenna with capacitively coupled parasitic elements. *IEEE Transactions on Antennas and Propagation*, *62*(2), 977-982.

[460] Oh, J., Thiel, M., and Sarabandi, K. (2014). Wave-propagation management in indoor environments using micro-radio-repeater systems. *IEEE Antennas and Propagation Magazine*, *56*(2), 76-88.

[461] Owen, K. J., VanDerElzen, B., Peterson, R. L., and Najafi, K. (2012, January). High aspect ratio deep silicon etching. In *Micro Electro Mechanical Systems (MEMS)* (pp. 251-254).

[462] Park, Y., and Wentzloff, D. D. (2010, May). A cyclic vernier time-to-digital converter synthesized from a 65nm CMOS standard library. In *Proceedings of 2010 IEEE International Symposium on Circuits and Systems* (pp. 3561-3564). IEEE.

[463] Park, Y., and Wentzloff, D. D. (2010, May). IR-UWB transmitters synthesized from standard digital library components. In *Proceedings of 2010 IEEE International Symposium on Circuits and Systems* (pp. 3296-3299). IEEE.

[464] Park, Y., and Wentzloff, D. D. (2010, September). An all-digital 12pJ/pulse 3.1–6.0 GHz IR-UWB transmitter in 65nm CMOS. In *2010 IEEE International Conference on Ultra-Wideband* (Vol. 1, pp. 1-4). IEEE.

[465] Park, Y., and Wentzloff, D. D. (2011). A cyclic Vernier TDC for ADPLLs synthesized from a standard cell library. *IEEE Transactions on Circuits and Systems I: Regular Papers*, *58*(7), 1511-1517.

[466] Park, Y., and Wentzloff, D. D. (2011). An all-digital 12 pJ/pulse IR-UWB transmitter synthesized from a standard cell library. *IEEE Journal of Solid-State Circuits*, *46*(5), 1147-1157.

[467] Park, Y., and Wentzloff, D. D. (2011, September). An all-digital PLL synthesized from a digital standard cell library in 65nm CMOS. In *2011 IEEE Custom Integrated Circuits Conference (CICC)* (pp. 1-4). IEEE.

[468] Peasley, B., Birchfield, S., Cunningham, A., and Dellaert, F. (2012, October). Accurate on-line 3D occupancy grids using Manhattan world constraints. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 5283-5290). IEEE.

[469] Peterson, K., and Fearing, R. S. (2011, September). Experimental dynamics of wing assisted running for a bipedal ornithopter. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on* (pp. 5080-5086). IEEE.

[470] Peterson, K., Birkmeyer, P., Dudley, R., and Fearing, R. S. (2011). A wing-assisted running robot and implications for avian flight evolution. *Bioinspiration & biomimetics*, *6*(4), 046008.

[471] Peterson, K., Birkmeyer, P., Dudley, R., and Fearing, R. S. (2011, March). Wing assisted locomotion of a 25 g running robot. In *Integrative and Comparative Biology*, *51*, E108-E108. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[472] Pippin, C., and Christensen, H. (2012, July). Performance based monitoring using statistical control charts on multi-robot teams. In

Information Fusion (FUSION), 2012 15th International Conference on (pp. 390-395). IEEE.

[473] Pullin, A. O., Kohut, N. J., Zarrouk, D., and Fearing, R. S. (2012, May). Dynamic turning of 13 cm robot comparing tail and differential drive. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on* (pp. 5086-5093). IEEE.

[474] Qian, F., Zhang, T., Li, C., Masarati, P., Hoover, A. M., Birkmeyer, P., and Goldman, D. I. (2012). Walking and running on yielding and fluidizing ground.

[475] Rais-Zadeh, M. (2012, May). Gallium nitride micromechanical resonators for IR detection. In *SPIE Defense, Security, and Sensing* (pp. 83731M-83731M). International Society for Optics and Photonics.

[476] Revzen, S., Burden, S. A., Moore, T. Y., Mongeau, J. M., and Full, R. J. (2013). Instantaneous kinematic phase reflects neuromechanical response to lateral perturbations of running cockroaches. *Biological cybernetics*, *107*(2), 179-200.

[477] Richter, C., Bry, A., and Roy, N. (2013, May). Polynomial trajectory planning for quadrotor flight. In *International Conference on Robotics and Automation*.

[478] Rinnerbauer, V., Ndao, S., Yeng, Y. X., Bermel, P., Senkevich, J., Celanovic, I., and Soljačić, M. (2012). High temperature nanophotonics: Controlling the flow of thermal radiation. In *Material Science-2012*, Chicago, IL.

[479] Rinnerbauer, V., Ndao, S., Yeng, Y. X., Chan, W. R., Senkevich, J. J., Joannopoulos, J. D., and Celanovic, I. (2012). Recent developments in high-temperature photonic crystals for energy conversion. *Energy & Environmental Science*, *5*(10), 8815-8823.

[480] Rinnerbauer, V., Ndao, S., Yeng, Y. X., Chan, W. R., Senkevich, J. J., Joannopoulos, J. D., and Celanovic, I. Recent developments in high-temperature photonic crystals. In *Royal Society of Chemistry: Energy and Environmental Science*, invited paper.

[481] Rinnerbauer, V., Senkevich, J. J., Joannopoulos, J. D., Soljačić, M., Celanovic, I., Harl, R. R., and Rogers, B. R. (2013). Low emissivity high-temperature tantalum thin film coatings for silicon devices. *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, *31*(1), 011501.

[482] Rinnerbauer, V., Yeng, Y. X., Chan, W. R., Senkevich, J. J., Joannopoulos, J. D., Soljačić, M., and Celanovic, I. (2012, September). Advanced high-temperature nanophotonics for efficient thermophotovoltaic energy conversion. In *27th EU PVSEC*, Frankfurt.

[483] Rogers III, J. G., Nieto-Granda, C., and Christensen, H. I. (2013). Coordination strategies for multi-robot exploration and mapping. In *Experimental Robotics* (pp. 231-243). Springer International Publishing.

[484] Rogers, J. G., Trevor, A. J., Nieto-Granda, C., and Christensen, H. I. (2011, September). Simultaneous localization and mapping with learned object recognition and semantic data association. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on* (pp. 1264-1270). IEEE.

[485] Sadeghi, M. M., Dowling, K., Peterson, R. L., and Najafi, K. (2013, January). High sensitivity, high density micro-hydraulic force sensor array utilizing stereo-lithography fabrication technique. In *Micro Electro Mechanical Systems (MEMS), 2013 IEEE 26th International Conference on* (pp. 673-676). IEEE.

[486] Sadeghi, M. M., Peterson, R. L., and Najafi, K. (2011, December). Micro-hydraulic structure for high performance bio-mimetic air flow sensor arrays. In *Electron Devices Meeting (IEDM), 2011 IEEE International* (pp. 29-4). IEEE.

[487] Sadeghi, M. M., Peterson, R. L., and Najafi, K. (2012). Hair-based sensors for micro-autonomous systems. In *Micro- and Nanotechnology Sensors, Systems, and Applications IV*, edited by T. George, M. Saif Islam, A. Dutta, Proceedings of the SPIE Vol. 8373 (SPIE, Bellingham, WA, 2012) 8373-1L.

[488] Sadeghi, M. M., Peterson, R. L., and Najafi, K. (2013, January). High-speed electrostatic micro-hydraulics for sensing and actuation. In *Micro Electro Mechanical Systems (MEMS), 2013 IEEE 26th International Conference on* (pp. 1191-1194). IEEE.

[489] Shah, S. K., and Tanner, H. G. (2013). Stochastic receding horizon control: Application to an octopedal robot. *Proc. SPIE, Micro-Nanotechnol. Sensors, Syst., Appl. V*, *8725*, 87251D-1.

[490] Shah, S. K., Pahlajani, C. D., Lacock, N. A., and Tanner, H. G. (2012, May). Stochastic receding horizon control for robots with probabilistic state constraints. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on* (pp. 2893-2898). IEEE.

[491] Shah, S. K., Pahlajani, C. D., Lacock, N. A., and Tanner, H. G. Receding horizon control for nonlinear stochastic systems with probabilistic state constraints. *IEEE Transactions on Automatic Control* (In preparation).

[492] Sharpe, S. S., Ding, Y., and Goldman, D. I. (2013). Environmental interaction influences muscle activation strategy during sand-swimming in the sandfish lizard Scincus scincus. *Journal of Experimental Biology*, *216*(2), 260-274.

[493] Shen, S., and Michael, N. (2013). State estimation for indoor and outdoor operation with a micro-aerial vehicle. In *Experimental Robotics* (pp. 273-288). Springer International Publishing.

[494] Shen, S., Michael, N., and Kumar, V. (2012). Stochastic differential equation-based exploration algorithm for autonomous indoor 3d exploration with a micro-aerial vehicle. *The International Journal of Robotics Research*, *31*(12), 1431-1444.

[495] Shen, S., Michael, N., and Kumar, V. (2012, May). Autonomous indoor 3D exploration with a micro-aerial vehicle. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on* (pp. 9-15). IEEE.

[496] Shrestha, E., Benedict, M., Hrishikeshavan, V., and Chopra, I. (2012). Development of a 100-gram micro-cyclocopter capable of autonomous hover.

[497] Sis, S. A., Lee, V., Phillips, J. D., and Mortazawi, A. (2012, June). A DC voltage dependent switchable acoustically coupled BAW filter based on BST-on-silicon composite structure. In *Microwave Symposium Digest (MTT), 2012 IEEE MTT-S International* (pp. 1-3). IEEE.

[498] Sis, S. A., Lee, V., Phillips, J. D., and Mortazawi, A. (2012, June). Intrinsically switchable thin film ferroelectric resonators. In *Microwave Symposium Digest (MTT), 2012 IEEE MTT-S International* (pp. 1-3). IEEE.

[499] Sis, S. A., Lee, V., Phillips, J. D., and Mortazawi, A. (2012, June). A DC voltage dependent switchable acoustically coupled BAW filter based on BST-on-silicon composite structure. In *Microwave Symposium Digest (MTT), 2012 IEEE MTT-S International* (pp. 1-3). IEEE.

[500] Song, Y. J., and Sarabandi, K. (2012). Equivalent Circuit Model for Metamaterial-Based Electromagnetic Band-Gap Isolator. *IEEE Antennas and Wireless Propagation Letters*, *11*, 1366-1369.

[501] Song, Y. J., and Sarabandi, K. (2012). Miniaturized radio repeater for enhanced wireless connectivity of ad-hoc networks. *IEEE Transactions on Antennas and Propagation*, *60*(8), 3913-3920.

[502] Song, Y. J., and Sarabandi, K. (2012, July). Design of radio repeater system using a near-field cancellation technique. In *Proceedings of the 2012 IEEE International Symposium on Antennas and Propagation* (pp. 1-2). IEEE.

[503] Springthorpe, D., Gravish, N., Mazouchova, N., Goldman, D., and Full, R. J. (2013, April). Burrowing biomechanics of the ghost crab. In *Integrative and Comparative Biology*, *53*, E206-E206. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[504] Sreetharan, P. S., Whitney, J. P., Strauss, M. D., and Wood, R. J. (2012). Monolithic manufacture of millimeter-scale machines. *J. Microengineering Micromechanics*, *22*(055027). (Cover article, IOP select article.)

[505] Sreetharan, P. S., Whitney, J. P., Strauss, M. D., and Wood, R. J. (2012). Monolithic fabrication of millimeter-scale machines. *Journal of Micromechanics and Microengineering*, *22*(5), 055027.

[506] Sui, Y., Gokhale, V. J., Shenderova, O. A., McGuire, G. E., and Rais-Zadeh, M. (2012). A thin-film infrared absorber using CNT/nanodiamond nanocomposite. In *MRS Proceedings* (Vol. 1452, pp. mrss12-1452). Cambridge University Press.

[507] Talucdher, A. R., Sergey, Y., Lingaiah, S., and Shivakumar, K. (2011). Indentation properties of NC dragonfly wing by nanolndentation. In *26th Annual Technical Conference of the American Society for Composites 2011 and the 2nd Joint US-Canada Conference on Composites*.

[508] Trevor, A. J., Rogers, J. G., and Christensen, H. I. (2012, May). Planar surface SLAM with 3D and 2D sensors. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on* (pp. 3041-3048). IEEE.

[509] Turpin, M., Michael, N., and Kumar, V. (2012). Trajectory design and control for aggressive formation flight with quadrotors. *Autonomous Robots*, *33*(1-2), 143-156.

[510] Turpin, M., Michael, N., and Kumar, V. (2013). Trajectory planning and assignment in multirobot systems. In *Algorithmic foundations of robotics X* (pp. 175-190). Springer Berlin

[511] Valbuena, L., and Tanner, H. G. (2012). Hybrid potential field based control of differential drive mobile robots. *Journal of intelligent & robotic systems*, 1-16.

[512] Venkata, S. K., and Jones, A. R. (2013). Leading-edge vortex structure over multiple revolutions of a rotating wing. *Journal of Aircraft*, *50*(4), 1312-1316.

[513] Wise, K. D. (2012, March). Wireless integrated microsystems: historical retrospective and future directions. In *Int. Conf. on Smart System Integration*, Zurich.

[514] Zarovy, S., Costello, M., and Mehta, A. (2013). Experimental method for studying gust effects on micro rotorcraft. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, *227*(4), 703-713.

[515] Zimmerman, J. D., Lee, K., and Forrest, S. R. (2012). *Integrated, flexible, high-efficiency solar cells: epitaxial lift-off GaAs solar cells and enabling substrate reuse*. Michigan Univ., Ann Arbor, MI.

FY13

[516] Adams, Z., Benedict, M., Hrishikeshavan, V., and Chopra, I. (2013). Design, development, and flight test of a small-scale cyclogyro UAV utilizing a novel cam-based passive blade pitching mechanism. *International Journal of Micro Air Vehicles*, *5*(2), 145-162.

[517] Benedict, M., Coleman, D., Mayo, D. B., and Chopra, I. (2013). Force and flowfield measurements on a rigid wing undergoing hover-capable flapping and pitching kinematics in air at MAV-scale Reynolds numbers. In *54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference* (p. 1706).

[518] Benedict, M., Mullins, J., Hrishikeshavan, V., and Chopra, I. (2013, January). Development of an optimized quad cycloidal-rotor UAV capable of autonomous stable hover. In *Proceedings of the American Helicopter Society International Specialists Meeting on Unmanned Rotorcraft, Scottsdale, AZ*.

[519] Benedict, M., Shrestha, E., Hrishikeshavan, V., and Chopra, I. (2014). Development of a micro twin-rotor cyclocopter capable of autonomous hover. *Journal of Aircraft*, *51*(2), 672-676.

[520] Bezzo, N., Anderson, M., Fierro, R., and Wood, J. (2012, November). Toward realistic coordination of connected heterogeneous robotic systems. In *In International Symposium on Distributed Autonomous Robotic Systems (DARS), Baltimore, MD*.

[521] Bezzo, N., Cortez, R. A., and Fierro, R. (2013). Exploiting Heterogeneity in Robotic Networks. In *Redundancy in Robot Manipulators and Multi-Robot Systems* (pp. 53-75). Springer Berlin Heidelberg.

[522] Bezzo, N., Griffin, B., Cruz, P., Donahue, J., Fierro, R., and Wood, J. (2014). A cooperative heterogeneous mobile wireless mechatronic system. *IEEE/ASME Transactions on Mechatronics*, *19*(1), 20-31.

[523] Bezzo, N., Sorrentino, F., and Fierro, R. (2013, June). Decentralized estimation of topology changes in wireless robotic networks. In *2013 American Control Conference* (pp. 5899-5904). IEEE.

[524] Buchan, A. D., Haldane, D. W., and Fearing, R. S. (2013, November). Automatic identification of dynamic piecewise affine models for a running robot. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on* (pp. 5600-5607). IEEE.

[525] Burden, S. A., and Sastry, S. S. (2013, June). Reduction and identification for hybrid dynamical models of terrestrial locomotion. In *SPIE Defense, Security, and Sensing* (pp. 87251B-87251B). International Society for Optics and Photonics.

[526] Burden, S. A., Revzen, S., Moore, T. Y., Sastry, S. S., and Full, R. J. (2013, April). Using reduced-order models to study dynamic legged locomotion: Parameter identification and model validation. In *Integrative and Comparative Biology*, *53*, E25-E25. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[527] Burden, S., Revzen, S., and Sastry, S. S. From anchors to templates: exact and approximate reduction in models of legged locomotion.

[528] Campbell, R., Humbert, J. S., and Chopra, I. (2013). Maneuverability comparison of flybar and flybarless micro-helicopters. In *AHS International Specialists' Meeting on Unmanned Rotorcraft and Network Centric Operations*, Scottsdale, AZ.

[529] Carpin, S., Chung, T. H., and Sadler, B. M. (2013, May). Theoretical foundations of high-speed robot team deployment. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on* (pp. 2033-2040). IEEE.

[530] Chan, W. R., Bermel, P., Pilawa-Podgurski, R. C., Marton, C. H., Jensen, K. F., Senkevich, J. J., ... & Celanovic, I. (2013). Toward high-energy-density, high-efficiency, and moderate-temperature chip-scale thermophotovoltaics. *Proceedings of the National Academy of Sciences*, *110*(14), 5309-5314.

[531] Chan, W. R., Bermel, P., Senkevich, J., Ghebrebrhan, M., Yeng, Y. X., Marton, C. H., Jensen, K. F., Johnson, S. G., Joannopoulos, J. D., Soljacic, M., and Celanovic, I. (2013). High energy density, high-efficiency, moderate-temperature chip-scale thermophotovoltaics. *Proceedings of the National Academy of Sciences*, *110*(14), 5309-5314.

[532] Choi, J., and Sarabandi, K. (2013, July). An extremely low-profile MF monopole antenna design and measurement. In *Radio Science Meeting (Joint with AP-S Symposium), 2013 USNC-URSI* (pp. 121-121). IEEE.

[533] Choi, J., Park, S., Cho, J., and Yoon, E. (2013, February). A 3.4μW CMOS image sensor with embedded feature-extraction algorithm for motion-triggered object-of-interest imaging. In *2013 IEEE International Solid-State Circuits Conference Digest of Technical Papers* (pp. 478-479). IEEE.

[534] Choi, J., Sheikhsofla, M., and Sarabandi, K. (2013, July). Near-ground wave propagation measurements at VHF ISM band. In *Radio Science Meeting (Joint with AP-S Symposium), 2013 USNC-URSI* (pp. 203-203). IEEE.

[535] Christensen, H. I., and Rogers, J. G. (2013, April). Robust multi-robot mapping and exploration. In *SPIE Defense, Security and Sensing*, Baltimore, MD.

[536] Conroy, J., Humbert, S. 2013, May 2). *Continuous-Time Structure from Motion using Omni-Directional Optic Flow for Micro Air Vehicle Applications.* SPIE-DSS, Baltimore, MD.

[537] Cortez, R. A., Fierro, R., and Wood, J. (2013). Connectivity maintenance of a heterogeneous sensor network. In *Distributed Autonomous Robotic Systems* (pp. 33-46). Springer Berlin Heidelberg.

[538] Cortez, R. A., Fierro, R., Wood, J., and Lumia, R. (2011, September). Heterogeneous sensor network for prioritized sensing. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 2333-2339). IEEE.

[539] Courtland, R. (2013). Better eyes for flying robots: new systems could improve the vision of micro aerial vehicles. In *IEEE Spectrum* online (http://spectrum.ieee.org/robotics/robotics-hardware/better-eyes-for-flying-robots), March 25, 2013.

[540] Courtland, R. (2013, April). Seeing on the fly. *IEEE Spectrum*, *50*(4), 14-16.

[541] Cruz, P., Fierro, R., Lu, W., and Ferrari, S. (2013, May). Maintaining robust connectivity in heterogeneous robotic networks. In *SPIE Defense, Security, and Sensing* (pp. 87410N-87410N). International Society for Optics and Photonics.

[542] Cunningham, A., Indelman, V., and Dellaert, F. (2013, May). DDF-SAM 2.0: Consistent distributed smoothing and mapping. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on* (pp. 5220-5227). IEEE.

[543] Cunningham, A., Ok, K., Antico, J., Indelman, V., and Dellaert, F. (2103, April). Aerial robot experimental design for decentralized visual SLAM. In *Unmanned Systems Technology XV, SPIE Defense, Security, and Sensing*, Baltimore, MD. SPIE.

[544] Desbiens, A. L., Pope, M., Berg, F., Teoh, Z. E., Lee, J., and Cutkosky, M. (2013, May). Efficient jumpgliding: theory and design considerations. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on* (pp. 4451-4458). IEEE.

[545] Ding, Y., Sharpe, S. S., Masse, A., and Goldman, D. I. (2012). Mechanics of undulatory swimming in a frictional fluid. *PLoS Comput Biol*, *8*(12), e1002810.

[546] Espinace, P., Kollar, T., Roy, N., and Soto, A. (2013). Indoor scene recognition by a mobile robot through adaptive object detection. *Robotics and Autonomous Systems*, *61*(9), 932-947.

[547] Faruque, I., Samuel, P. D., and Humbert, J. (2014). Yaw dynamics identification of an insect-inspired flapping wing micro air vehicle. In *AIAA Guidance, Navigation, and Control Conference* (p. 1468).

[548] Faust, A., Palunko, I., Cruz, P., Fierro, R., and Tapia, L. (2013, May). Learning swing-free trajectories for UAVs with a suspended load. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on* (pp. 4902-4909). IEEE.

[549] Fink, J., Twigg, J., Yu, P., and Sadler, B. M. (2013, December). A parsimonious model for wireless connectivity in robotic networks. In *Global Conference on Signal and Information Processing (GlobalSIP), 2013 IEEE* (pp. 855-858). IEEE.

[550] Fu, J., Shah, S., and Tanner, H. G. (2013). Hierarchical control via simulation relations and feedback linearization. In *American Control Conference*.

[551] Gardner, R. C., and Humbert, J. S. (2014). Comparative framework for maneuverability and gust tolerance of microhelicopters. *Journal of Aircraft*, *51*(5), 1546-1553.

[552] Ghosh, D., and Baeder, J. D. (2012). Compact reconstruction schemes with weighted ENO limiting for hyperbolic conservation laws. *SIAM Journal on Scientific Computing*, *34*(3), A1678-A1706.

[553] Ghosh, D., Medida, S., and Baeder, J. (2012, November). Direct numerical simulation of compressible turbulent flows with weighted non-linear compact schemes. In *APS Division of Fluid Dynamics Meeting Abstracts* (Vol. 1, p. 5006).

[554] Gokhale, V. J., Shenderova, O. A., McGuire, G. E., and Rais-Zadeh, M. (2014). Infrared absorption properties of carbon nanotube/nanodiamond based thin film coatings. *Journal of Microelectromechanical Systems*, *23*(1), 191-197.

[555] Gravish, N., Monaenkova, D., Goodisman, M. A., and Goldman, D. I. (2013). Climbing, falling, and jamming during ant locomotion in confined environments. *Proceedings of the National Academy of Sciences*, *110*(24), 9746-9751.

[556] Haldane, D. W., Peterson, K. C., Bermudez, F. L. G., and Fearing, R. S. (2013, May). Animal-inspired design and aerodynamic stabilization of a hexapedal millirobot. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on* (pp. 3279-3286). IEEE.

[557] Harrington, A. M., and Kroninger, C. (2013). *Characterization of small DC brushed and brushless motors* (No. ARL-TR-6389). US Army Research Laboratory, Vehicle Technology Directorate, Aberdeen Proving Ground, MD.

[558] Hatton, R. L., Ding, Y., Choset, H., and Goldman, D. I. (2013). Geometric visualization of self-propulsion in a complex medium. *Physical review letters*, *110*(7), 078101.

[559] Hawkes, E. W., Eason, E. V., Asbeck, A. T., and Cutkosky, M. R. (2013). The gecko's toe: Scaling directional adhesives for climbing applications. *IEEE/ASME Transactions on Mechatronics*, *18*(2), 518-526.

[560] Hawkes, E. W., Eason, E. V., Jiang, H., Estrada, M. A., Pope, M. T., Christensen, D. L., Cutkosky, M. R., Kehlenbeck, A., and Humbert, J. S. (2013). Dynamic surface grasping for UAVs. In *RSS 2013 Workshop on Aerial Mobile Manipulation*, Berlin, 27 June 2013.

[561] Helicopter Society Specialists' Meeting for Unmanned Rotorcraft, Scottsdale, AZ, Jan. 2013.

[562] Hrishikeshavan, V, Benedict, M., and Chopra, I. (2013). Flight dynamics system identification and control of a cyclocopter micro air vehicle in hover. In *Proceedings of the 69th Annual National Forum of the American Helicopter Society*, Phoenix, AZ, May 21–23, 2013.

[563] Hrishikeshavan, V., Humbert, J. S., and Chopra, I. (2014). Gramian Analysis of a Shrouded Rotor Micro Air Vehicle in Hover. *Journal of Guidance, Control, and Dynamics*, *37*(5), 1684-1691.

[564] Hrishikeshavan, V., Humbert, J.S., and Chopra, I. (2014). Controllability and disturbance Gramian of a shrouded rotor micro air vehicle. *AIAA Journal of Guidance, Control and Dynamics*, *37*(5), 1684-1691.

[565] Jarugumilli, T., Benedict, M., Lind, A. H., and Chopra, I. (2013). Performance and flow visualization studies to examine the role of pitching kinematics on MAV-scale cycloidal rotor performance in forward flight. In *Proceedings of the American Helicopter Society International Specialists' Meeting on Unmanned Rotorcraft*, Scottsdale, AZ, January 22-24, 2013.

[566] Jarugumilli, T., Lind, A., Benedict, M., Lakshminarayan, V., Jones, A., and Chopra, I. (2013, May). Experimental and computational flow field studies of a MAV-scale cycloidal rotor in forward flight. In *69th Annual International Forum, American Helicopter Society*.

[567] Jeon, D., Kim, Y., Lee, I., Zhang, Z., Blaauw, D., and Sylvester, D. (2013, February). A 470mV 2.7 mW feature extraction-accelerator for micro-autonomous vehicle navigation in 28nm CMOS. In *2013 IEEE International Solid-State Circuits Conference Digest of Technical Papers* (pp. 166-167). IEEE.

[568] Kohut, N. J., Pullin, A. O., Haldane, D. W., Zarrouk, D., and Fearing, R. S. (2013, May). Precise dynamic turning of a 10 cm legged robot on a low friction surface using a tail. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on* (pp. 3299-3306). IEEE.

[569] Kohut, N. J., Zarrouk, D., Peterson, K. C., and Fearing, R. S. (2013, November). Aerodynamic steering of a 10 cm high-speed running robot. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on* (pp. 5593-5599). IEEE.

[570] Kolluru Venkata, S., and Jones, A. R. (2013). Leading edge vortex structure over multiple revolutions of a rotating wing. *AIAA Journal, 50*(4), 1312-1316.

[571] Lakshminarayan, V. K., and Baeder, J. D. (2012). Improved shroud design for rotary wing MAV applications based on computational analysis. *Journal of the American Helicopter Society*, *57*(4), 1-5.

[572] Lee, K., Zimmerman, J. D., Hughes, T. W., and Forrest, S. R. (2014). Non-destructive wafer recycling for low-cost thin-film flexible optoelectronics. *Advanced Functional Materials*, *24*(27), 4284-4291.

[573] Lee, S., Lee, V., Sis, S. A., and Mortazawi, A. (2013). Large-signal performance and modeling of intrinsically switchable ferroelectric FBARs. *IEEE Transactions on Microwave Theory and Techniques*, *61*(1), 415-422.

[574] Lee, S., Lee, V., Sis, S. A., and Mortazawi, A. (2013, June). Linearity analysis of intrinsically switchable ferroelectric FBAR filters. In *Microwave Symposium Digest (IMS), 2013 IEEE MTT-S International* (pp. 1-3). IEEE.

[575] Lee, V., Sis, S. A., Phillips, J. D., and Mortazawi, A. (2013). Intrinsically Switchable Ferroelectric Contour Mode Resonators. *IEEE Transactions on Microwave Theory and Techniques*, *61*(8), 2806-2813.

[576] Lee, V.; Lee, S.; Sis, S.A.; Mortazawi, A. (2013). Reliability of intrinsically switchable ferroelectric resonators and filters. In *Microwave Symposium Digest (MTT), 2013 IEEE MTT-S International* (pp.1-3), 2-7 June 2013. IEEE.

[577] Li, C., Zhang, T., and Goldman, D. I. (2013). A terradynamics of legged locomotion on granular media. *Science*, *339*(6126), 1408-1412.

[578] Lin, D. T., Chae, H., Li, L., and Flynn, M. P. (2013). A low-power adaptive receiver utilizing discrete-time spectrum-sensing. *IEEE Transactions on Microwave Theory and Techniques*, *61*(3), 1338-1346.

[579] Luna, J. M., Fierro, R., Abdallah, C. T., and Lewis, F. (2012, December). A binary consensus approach to decentralized coordination of nonholonomic sensor networks. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)* (pp. 3353-3359). IEEE.

[580] Luna, J. M., Fierro, R., Abdallah, C. T., and Lewis, F. (2013). On the compression of locational and environmental data in multi-vehicle missions: a control systems approach. *International Journal of Control*, *86*(7), 1281-1291.

[581] Luna, J. M., Fierro, R., Abdallah, C. T., and Wood, J. (2013). An adaptive coverage control for deployment of nonholonomic mobile sensor networks over time-varying sensory functions. *Asian Journal of Control*, *15*(4), 988-1000.

[582] MacFarlane, K., Kostreski, N., and Humbert, J.S. (2013). Forward flight dynamics of an insect-inspired flapping wing vehicle. In *AHS International Specialists Meeting on Unmanned Rotorcraft and Network Centric Operations*. January 22-24, 2013.

[583] Malhan, R., Baeder, J., Chopra, I., and Masarati, P. (2013). Investigation of Aerodynamics of Flapping Wings for MAV Applications Using Experiments and CFD–CSD Analysis. In *Proceedings of the American Helicopter Society 5th International Specialists Meeting on Unmanned Rotorcraft and Network Centric Operations, Scottsdale, AZ*.

[584] Malhan, R., Benedict, M., and Chopra, I. (2012). Experimental studies to understand the hover and forward flight performance of a MAV-scale flapping wing concept. *Journal of the American Helicopter Society*, *57*(2), 1-11.

[585] Mayo, D., and Jones, A. (2013, January). Evolution and breakdown of a leading edge vortex on a rotating wing. In *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition* (p. 843).

[586] Mayo, D., and Jones, A. R. (2013).The effect of wake encounters on the leading edge vortex generated by a rotating wing. In *the 51st AIAA Aerospace Sciences Meeting and Exhibit*, Grapevine, TX, Jan. 7-10, 2013.

[587] Mazouchova, N., Umbanhowar, P. B., and Goldman, D. I. (2013). Flipper-driven terrestrial locomotion of a sea turtle-inspired robot. *Bioinspiration & biomimetics*, *8*(2), 026007.

[588] Milligan, J., Hsieh, M. A., and Chaimowicz, L. (2013, August). Vision-Based Cooperative Localization for Small Networked Robot Teams. In *Conference Towards Autonomous Robotic Systems* (pp. 161-172). Springer Berlin Heidelberg.

[589] Milligan, T. J. (2013). *Collaborative localization for small networked robot teams* (Doctoral dissertation, Drexel University).

[590] Moallem, M., and Sarabandi, K. (2013, January). Y-band radar assessment for indoor navigation and mapping. In *Radio Science Meeting (USNC-URSI NRSM), 2013 US National Committee of URSI National* (pp. 1-1). IEEE.

[591] Mongeau, J. M., Demir, A., Lee, J., Cowan, N. J., and Full, R. J. (2013). Locomotion-and mechanics-mediated tactile sensing: antenna

reconfiguration simplifies control during high-speed navigation in cockroaches. *Journal of Experimental Biology*, *216*(24), 4530-4541.

[592] Mukherjee, R. (2013, August). Parallel algorithm for constrained multi-rigid body system dynamics in generalized topologies. In *ASME 2013 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (pp. V07AT10A036-V07AT10A036). American Society of Mechanical Engineers.

[593] Mukherjee, R., and Laflin, J. (2013, August). Parallel algorithm for modeling constrained multi-flexible body system dynamics. In *ASME 2013 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (pp. V07BT10A010-V07BT10A010). American Society of Mechanical Engineers.

[594] Mukherjee, R., and Malczyk, P. (2013, August). Efficient approach for constraint enforcement in constrained multibody system dynamics. In *ASME 2013 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (pp. V07AT10A037-V07AT10A037). American Society of Mechanical Engineers.

[595] Mukherjee, R., and Malczyk, P. (2013, August). Parallel algorithm for modeling multi-rigid body system dynamics with nonholonomic constraints. In *ASME 2013 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (pp. V07AT10A038-V07AT10A038). American Society of Mechanical Engineers.

[596] Najafi, K. (2012, October). Biomimetic hair sensors: utilizing the third dimension. In *Sensors, 2012 IEEE* (pp. 1-4). IEEE.

[597] Najafi, K. (2013, December). 3-D biomimetic hair sensors and actuators. In *2013 IEEE International Electron Devices Meeting* (pp. 8-2). IEEE.

[598] Nieto-Granda, C., Rogers III, J. G., and Christensen, H. I. (2013, May). Multi-robot exploration strategies for tactical tasks in urban environments. In *SPIE Defense, Security, and Sensing* (pp. 87410B-87410B). International Society for Optics and Photonics.

[599] Nieto-Granda, C., Rogers III, J. G., and Christensen, H. I. (2013, May). Multi-robot exploration strategies for tactical tasks in urban environments. In *SPIE Defense, Security, and Sensing* (pp. 87410B-87410B). International Society for Optics and Photonics.

[600] Oh, J., and Sarabandi, K. (2013). A topology-based miniaturization of circularly polarized patch antennas. *IEEE Transactions on Antennas and Propagation*, *61*(3), 1422-1426.

[601] Oh, J., and Sarabandi, K. (2014). Compact, low profile, common aperture polarization, and pattern diversity antennas. *IEEE Transactions on Antennas and Propagation*, *62*(2), 569-576.

[602] Oh, J., Choi, J., Dagefu, F. T., and Sarabandi, K. (2013). Extremely small two-element monopole antenna for HF band applications. *IEEE Transactions on Antennas and Propagation*, *61*(6), 2991-2999.

[603] Palunko, I., Faust, A., Cruz, P., Tapia, L., and Fierro, R. A reinforcement learning approach to suspended load manipulation with aerial robots. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)* (pp. 4887-4894).

[604] Pivtoraiko, M. (2013, May). Generating fast maneuvers for groups of networked UAVs. At the *Workshop on Networked Multi-Agent Systems: From Theory to Practice*, collocated with the *IEEE International Conference on Robotics and Automation*.

[605] Pivtoraiko, M., Mellinger, D., and Kumar, V. (2013, May). Incremental micro-UAV motion replanning for exploring unknown environments. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on* (pp. 2452-2458). IEEE.

[606] Qawasmeh, B. R., and Wei, M. (2013). Low-dimensional models for compressible temporally developing shear layers. *Journal of Fluid Mechanics*, *731*, 364-393.

[607] Qian, F. , Daffon, K., Zhang, T., and Goldman, D. I. (2013). An automated system for systematic testing of locomotion on heterogeneous granular media. In *Proceedings of the 16th International Conference on Climbing and Walking Robots (CLAWAR)* (pp. 1-8).

[608] Rahimi, N., Aragon, A. A., Romero, O. S., Shima, D. M., Rotter, T. J., Balakrishnan, G., and Lester, L. F. (2013, June). Ultra-low resistance NiGeAu and PdGeAu ohmic contacts on N-GaSb grown on GaAs. In *Photovoltaic Specialists Conference (PVSC), 2013 IEEE 39th* (pp. 2123-2126). IEEE.

[609] Revzen, S., Burden, S., Koditschek, D., and Sastry, S. S. (2013). Pinned equilibria provide robustly stable multilegged locomotion. *Dynamic Walking*.

[610] Rinnerbauer, V., Ndao, S., Xiang Yeng, Y., Senkevich, J. J., Jensen, K. F., Joannopoulos, J. D., ... & Geil, R. D. (2013). Large-area fabrication of high aspect ratio tantalum photonic crystals for high-temperature selective emitters. *Journal of Vacuum Science & Technology B, Nanotechnology and Microelectronics: Materials, Processing, Measurement, and Phenomena*, *31*(1), 011802.

[611] Rinnerbauer, V., Senkevich, J. J., Joannopoulos, J. D., Soljačić, M., Celanovic, I., Harl, R. R., and Rogers, B. R. (2013). Low emissivity high-temperature tantalum thin film coatings for silicon devices. *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, *31*(1), 011501.

[612] Rinnerbauer, V., Yeng, Y. X., Chan, W. R., Senkevich, J. J., Joannopoulos, J. D., Soljačić, M., and Celanovic, I. (2013). High-temperature stability and selective thermal emission of polycrystalline tantalum photonic crystals. *Optics express*, *21*(9), 11482-11491.

[613] Roberts, N. E., and Wentzloff, D. D. (2012, October). 915MHz ultra low power receiver using sub-V t active rectifiers. In *Subthreshold Microelectronics Conference (SubVT), 2012 IEEE* (pp. 1-3). IEEE.

[614] Rogers III, J. G., Nieto-Granda, C., and Christensen, H. I. (2013). Coordination strategies for multi-robot exploration and mapping. In *Experimental Robotics* (pp. 231-243). Springer International Publishing.

[615] Rogers, J. G., and Christensen, H. I. (2013, May). Robot planning with a semantic map. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on* (pp. 2239-2244). IEEE.

[616] Rogers, J. G., Young, S., Gregory, J., Nieto-Granda, C., and Christensen, H. I. (2013, May). Robot mapping in large-scale mixed indoor and outdoor environments. In *SPIE Defense, Security, and Sensing* (pp. 874107-874107). International Society for Optics and Photonics.

[617] Sadeghi, M. M., Dowling, K., Peterson, R. L., and Najafi, K. (2013, January). High sensitivity, high density micro-hydraulic force sensor array utilizing stereo-lithography fabrication technique. In *Micro Electro Mechanical Systems (MEMS), 2013 IEEE 26th International Conference on* (pp. 673-676). IEEE.

[618] Sadeghi, M. M., Peterson, R. L., and Najafi, K. (2013, January). High-speed electrostatic micro-hydraulics for sensing and actuation. In *Micro Electro Mechanical Systems (MEMS), 2013 IEEE 26th International Conference on* (pp. 1191-1194). IEEE.

[619] Sadeghi, M. M., Peterson, R. L., and Najafi, K. (2013, June). A 2-D directional air flow sensor array made using stereolithography and MEMS micro-hydraulic structures. In *2013 Transducers & Eurosensors XXVII: The 17th International Conference on Solid-State Sensors, Actuators and Microsystems (Transducers & Eurosensors XXVII)* (pp. 722-725). IEEE.

[620] Schlueter, K. L., Jones, A. R., Granlund, K., and Ol, M. (2014). Effect of root cutout on force coefficients of rotating wings. *AIAA journal*, *52*(6), 1322-1325.

[621] Schlueter, K., Jones, A. R., Granlund, K., and Ol, M. Effect of root cutout on the lift produced by a rotating wing. To appear in *AIAA Journal*.

[622] Sharpe, S. S., Ding, Y., and Goldman, D. I. (2013). Environmental interaction influences muscle activation strategy during sand-swimming in the sandfish lizard Scincus scincus. *Journal of Experimental Biology*, *216*(2), 260-274.

[623] Sheikhsofla, M., and Sarabandi, K. (2013, July). Indoor wave propagation simulations at HF using rayleigh-gans approximation. In *Radio Science Meeting (Joint with AP-S Symposium), 2013 USNC-URSI* (pp. 202-202). IEEE.

[624] Shen, S., Mulgaonkar, Y., Michael, N., and Kumar, V. (2013, May). Vision-based state estimation for autonomous rotorcraft mavs in complex environments. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on* (pp. 1758-1764). IEEE.

[625] Song, Y. J., and Sarabandi, K. (2012). Equivalent circuit model for metamaterial-based electromagnetic band-gap isolator. *IEEE Antennas and Wireless Propagation Letters*, *11*, 1366-1369.

[626] Song, Y. J., and Sarabandi, K. (2013, July). Gain and bandwidth enhancement of micro-radio-repeater for ad-hoc wireless communication. In *IEEE International Symposium on Antenna and Propagation Society*.

[627] Song, Y. J., and Sarabandi, K. (2014). A simultaneous dual-channel micro-radio-repeater for ad-hoc wireless communication. *IEEE Transactions on Antennas and Propagation*, *62*(6), 3378-3383.

[628] Stelmakh, V., Rinnerbauer, V., Geil, R. D., Aimone, P. R., Senkevich, J. J., Joannopoulos, J. D., and Celanovic, I. (2013). High-temperature tantalum tungsten alloy photonic crystals: Stability, optical properties, and fabrication. *Applied Physics Letters*, *103*(12), 123903.

[629] Stelmakh, V., Rinnerbauer, V., Joannopoulos, J. D., Soljačić, M., Celanovic, I., Senkevich, J. J., and Shrader, R. (2013). Evolution of sputtered tungsten coatings at high temperature. *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, *31*(6), 061505.

[630] Tang, Y., Peterson, R. L., and Najafi, K. (2013, January). Technology for fabricating dense 3-D microstructure arrays for biomimetic hair-like sensors. In *Micro Electro Mechanical Systems (MEMS), 2013 IEEE 26th International Conference on* (pp. 355-358). IEEE.

[631] Thomas, J., Polin, J., Sreenath, K., and Kumar, V. (2013, August). Avian-inspired grasping for quadrotor micro uavs. In *ASME 2013 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (pp. V06AT07A014-V06AT07A014). American Society of Mechanical Engineers.

[632] Thomas, J., Pope, M., Loianno, G., Hawkes, E. W., Estrada, M. A., Jiang, H., and Kumar, V. (2016). Aggressive flight with quadrotors for perching on inclined surfaces. *Journal of Mechanisms and Robotics*, *8*(5), 051007.

[633] Twigg, J. N., Fink, J., Paul, L. Y., and Sadler, B. M. (2013). Efficient base station connectivity area discovery. *The International Journal of Robotics Research*, 0278364913488634.

[634] Venkata, S. K., and Jones, A. R. (2013). Leading-edge vortex structure over multiple revolutions of a rotating wing. *Journal of Aircraft*, *50*(4), 1312-1316.

[635] Venkateswaran, S., Isaacs, J. T., Fregene, K., Ratmansky, R., Sadler, B. M., Hespanha, J. P., and Madhow, U. (2013, June). RF source-seeking by a micro aerial vehicle using rotation-based angle of arrival estimates. In *2013 American Control Conference* (pp. 2581-2587). IEEE.

[636] Wei, M., Xu, M., and Yang, T. (2012, November). Global model reduction for fluid-structure system. *APS Meeting Abstracts, 1*, 25006.

[637] Weiss, S., Brockers, R., and Matthies, L. (2013, November). 4DoF drift free navigation using inertial cues and optical flow. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 4180-4186). IEEE.

[638] Wu, X. A., Suresh, S. A., Jiang, H., Ulmen, J. V., Hawkes, E. W., Christensen, D. L., and Cutkosky, M. R. (2015, September). Tactile sensing for gecko-inspired adhesion. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on* (pp. 1501-1507). IEEE.

[639] Yan, Y., and Mostofi, Y. (2013, June). Communication and path planning strategies of a robotic coverage operation. In *2013 American Control Conference* (pp. 860-866). IEEE.

[640] Yang, T., and Wei, M. (2013). A fully-coupled approach to simulate three-dimensional flexible flapping wings. In *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition* (p. 864).

[641] Yeng, Y. X., Chan, W. R., Rinnerbauer, V., Joannopoulos, J. D., Soljačić, M., and Celanovic, I. (2013). Performance analysis of experimentally viable photonic crystal enhanced thermophotovoltaic systems. *Optics express*, *21*(106), A1035-A1051.

[642] Zachary, H., A., Benedict, M., Hrishikeshavan, V., and Chopra, I., "Design and Development of a Small-Scale Cyclogyro UAV Utilizing a Novel Cam-Based Passive Blade Pitching Mechanism," for publication in the International Journal of Micro Air Vehicles.

[643] Zarrouk, D., Pullin, A. O., Kohut, N. J., and Fearing, R. S. (2013). Sprawl tuned dynamic hexapod. In *IEEE Int. Conf. on Robotics and Automation*, Karlsruhe, Germany May 6-10, 2013.

[644] Zhang, T., Qian, F., Li, C., Masarati, P., Hoover, A. M., Birkmeyer, P., and Goldman, D. I. (2013). Ground fluidization promotes rapid running of a lightweight robot. *The International Journal of Robotics Research*, *32*(7), 859-869.

FY14

[645] Adams, Z., Benedict, M., Hrishikeshavan, V., and Chopra, I. (2013). Design, development, and flight test of a small-scale cyclogyro UAV utilizing a novel cam-based passive blade pitching mechanism. *International Journal of Micro Air Vehicles*, *5*(2), 145-162.

[646] Ansari, A., and Rais-Zadeh, M. (2014). A temperature-compensated micromechanical gallium nitride resonator. *IEEE Electron Device Letters*, *35*(11), 1127-1129.

[647] Ansari, A., and Rais-Zadeh, M. (2014, June). An 8.7 GHz GaN micromechanical resonator with an integrated AlGaN/GaN HEMT. In *Proceedings of the Solid-State Sensors, Actuators and Microsystems Workshop (Hilton Head Workshop), Hilton Head, CA, USA* (Vol. 812, p. 295296).

[648] Badrya, C., Amiraux, M., Baeder, J. D., and Kroninger, C. M. (2014, January 22-24). Parametric study of insect flapping performance using CFD. Badrya, C., and Baeder, J., *AHS Aeromechanics Specialists' Conference*, San Francisco, CA.

[649] Badrya, C., Baeder, J.D., Amiraux, M., and Kroninger, C. (2014.) Parametric study of insect flapping performance using CFD. In *The Fifth Decennial American Helicopter Society Specialists' Conference*, San Francisco, California, Jan 22– 24, 2014.

[650] Beals, N., and Jones, A. R. (2014, June). Unsteady lift generation on a passively deforming rotating wing revisited. AIAA Paper 2014-2138 presented at the *32nd AIAA Applied Aerodynamics Conference*, Atlanta, GA.

[651] Benedict, M., Jarugumilli, T., Lakshminarayan, V., and Chopra, I. (2014). Effect of flow curvature on forward flight performance of a micro-air-vehicle-scale cycloidal-rotor. *AIAA Journal*, *52*(6), 1159-1169.

[652] Benedict, M., Shrestha, E., Hrishikeshavan, V., and Chopra, I. (2014). Development of a micro twin-rotor cyclocopter capable of autonomous hover. *Journal of Aircraft*, *51*(2), 672-676.

[653] Bezzo, N., Anderson, M. S., Fierro, R., and Wood, J. (2014). A real world coordination framework for connected heterogeneous robotic systems. In *Distributed Autonomous Robotic Systems* (pp. 75-89). Springer Berlin Heidelberg.

[654] Bezzo, N., Cruz, P. J., Sorrentino, F., and Fierro, R. (2014). Decentralized identification and control of networks of coupled mobile platforms through adaptive synchronization of chaos. *Physica D: Nonlinear Phenomena*, *267*, 94-103.

[655] Bezzo, N., Griffin, B., Cruz, P., Donahue, J., Fierro, R., and Wood, J. (2014). A cooperative heterogeneous mobile wireless mechatronic system. *IEEE/ASME Transactions on Mechatronics*, *19*(1), 20-31.

[656] Brockers, R., Kuwata, Y., Weiss, S., and Matthies, L. (2014, June). Micro air vehicle autonomous obstacle avoidance from stereo-vision. In *SPIE Defense+ Security* (pp. 90840O-90840O). International Society for Optics and Photonics.

[657] Buchan, A. D., Haldane, D. W., and Fearing, R. S. (2013, November). Automatic identification of dynamic piecewise affine models for a running robot. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 5600-5607). IEEE.

[658] Burden, S. A., Gonzalez, H., Vasudevan, R., Bajcsy, R., and Sastry, S. S. (2015). Metrization and simulation of controlled hybrid systems. *IEEE Transactions on Automatic Control*, *60*(9), 2307-2320.

[659] Burden, S. A., Revzen, S., and Sastry, S. S. (2015). Model reduction near periodic orbits of hybrid dynamical systems. *IEEE Transactions on Automatic Control*, *60*(10), 2626-2639.

[660] Burden, S. A., Sastry, S. S., and Full, R. J. (2014, January). Optimization for models of legged locomotion: parameter estimation, gait synthesis, and experiment design. In *INTEGRATIVE AND COMPARATIVE BIOLOGY* (Vol. 54, pp. E28-E28). JOURNALS DEPT, 2001 EVANS RD, CARY, NC 27513 USA: OXFORD UNIV PRESS INC.

[661] Carpin, S., Pavone, M., and Sadler, B. M. (2014, September). Rapid multirobot deployment with time constraints. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 1147-1154). IEEE.

[662] Chan, W. R., Waits, C. M., Joannopoulos, J. D., and Celanovic, I. (2014, June). Thermophotovoltaic and thermoelectric portable power generators. In *SPIE Defense+ Security* (pp. 90831W-90831W). International Society for Optics and Photonics.

[663] Chen Li, Ronald Fearing, Robert Full. (2014). Biomechanics of insects traversing vertical compliant beam obstacles. World Congress of Biomechanics. Boston, MA.

[664] Cheng, M. X., Ling, Y., and Sadler, B. M. (2014, December). Wireless ad hoc networks connectivity assessment and relay node deployment. In *2014 IEEE Global Communications Conference* (pp. 399-404). IEEE.

[665] Cho, J., Park, S., Choi, J., and Yoon, E. (2014, July). Area-efficient and low-power implementation of vision chips using multi-level mixed-mode processing. In *2014 14th International Workshop on Cellular Nanoscale Networks and their Applications (CNNA)* (pp. 1-2). IEEE.

[666] Choi, J., and Sarabandi, K. (2013, July). An extremely low-profile MF monopole antenna design and measurement. In *Radio Science Meeting (Joint with AP-S Symposium), 2013 USNC-URSI* (pp. 121-121). IEEE.

[667] Choi, J., and Sarabandi, K. (2014, July). Highly miniaturized low-VHF folded dipole antenna for compact, mobile communication applications. In *Radio Science Meeting (Joint with AP-S Symposium), 2014 USNC-URSI* (pp. 30-30). IEEE.

[668] Choi, J., Park, S., Cho, J., and Yoon, E. (2014). A 3.4-W object-adaptive CMOS image sensor with embedded feature extraction algorithm for motion-triggered object-of-interest imaging. *IEEE Journal of Solid-State Circuits*, *49*(1), 289-300.

[669] Choi, J., Sheikhsofla, M., and Sarabandi, K. (2013, July). Near-ground wave propagation measurements at VHF ISM band. In *Radio Science Meeting (Joint with AP-S Symposium), 2013 USNC-URSI* (pp. 203-203). IEEE.

[670] Chou, J. B., Yeng, Y. X., Lee, Y. E., Lenert, A., Rinnerbauer, V., Celanovic, I., ... & Kim, S. G. (2014). Enabling ideal selective solar absorption with 2D metallic dielectric photonic crystals. *Advanced Materials*, *26*(47), 8041-8045.

[671] Chou, J. B., Yeng, Y. X., Lenert, A., Rinnerbauer, V., Celanovic, I., Soljačić, M., and Kim, S. G. (2014). Design of wide-angle selective absorbers/emitters with dielectric filled metallic photonic crystals for energy applications. *Optics express*, *22*(101), A144-A154.

[672] Choudhary, S., Trevor, A. J., Christensen, H. I., and Dellaert, F. (2014, September). SLAM with object discovery, modeling and mapping. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 1018-1025). IEEE.

[673] Conroy, J., Kehlenbeck, A., Humbert, J. S., and Nothwang, W. (2014, June). Characterization and enhancement of micro brushless DC motor

response. In *SPIE Defense+ Security* (pp. 90831P-90831P). International Society for Optics and Photonics.

[674] Conroy, J., Kehlenbeck, A., Humbert, J. S., and Nothwang, W. (2014, June). Characterization and enhancement of micro brushless DC motor response. In *SPIE Defense+ Security* (pp. 90831P-90831P). International Society for Optics and Photonics.

[675] Cruz, P., and Fierro, R. (2014, October). Autonomous lift of a cable-suspended load by an unmanned aerial robot. In *2014 IEEE Conference on Control Applications (CCA)* (pp. 802-807). IEEE.

[676] Cruz, P., Valbuena, L., and Fierro, R. (2014). Synchronization of heterogeneous wireless robotic networks. In *The 2nd Annual Workshop on Wireless Intelligent Sensor Networks (WISeNET)*, Durham, NC, June 9-10, 2014.

[677] Dagefu, F. T., and Sarabandi, K. (2010, July). Simulation and measurement of near-ground wave propagation for indoor scenarios. In *Antennas and Propagation Society International Symposium (APSURSI), 2010 IEEE* (pp. 1-4). IEEE.

[678] Dagefu, F. T., Oh, J., and Sarabandi, K. (2013). A sub-wavelength rf source tracking system for GPS-denied environments. *IEEE Transactions on Antennas and Propagation*, *61*(4), 2252-2262.

[679] Dagefu, F. T., Oh, J., Choi, J., and Sarabandi, K. (2013). Measurements and physics-based analysis of co-located antenna pattern diversity system. *IEEE Transactions on Antennas and Propagation*, *61*(11), 5724-5734.

[680] Dagefu, F. T., Verma, G., Rao, C. R., Paul, L. Y., Fink, J. R., Sadler, B. M., and Sarabandi, K. (2015). Short-range Low-VHF channel characterization in cluttered environments. *IEEE Transactions on Antennas and Propagation*, *63*(6), 2719-2727.

[681] Desaraju, V. R., and Michael, N. (2014, May). Hierarchical adaptive planning in environments with uncertain, spatially-varying disturbance forces. In *2014 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 5171-5176). IEEE.

[682] Desaraju, V. R., Liu, L., and Michael, N. (2016). Multi-vehicle adaptive planning with online estimated cost due to disturbance forces. In *Intelligent Autonomous Systems 13* (pp. 321-334). Springer International Publishing.

[683] Desaraju, V. R., Michael, N., Humenberger, M., Brockers, R., Weiss, S., and Matthies, L. (2014). Vision-based landing site evaluation and trajectory generation toward rooftop landing. In *Proceedings of robotics: science and systems*.

[684] Desbiens, A. L., Pope, M. T., Christensen, D. L., Hawkes, E. W., and Cutkosky, M. R. (2014). Design principles for efficient, repeated

jumpgliding. *Bioinspiration & biomimetics*, *9*(2), 025009. This paper includes material presented in [1] with major extensions.

[685] Ding, Y., Li, C., and Goldman, D. I. (2013). Swimming in the desert. *Physics Today*, *66*(11), 68.

[686] Ding, Y., Sharpe, S. S., Wiesenfeld, K., and Goldman, D. I. (2013). Emergence of the advancing neuromechanical phase in a resistive force dominated medium. *Proceedings of the National Academy of Sciences*, *110*(25), 10123-10128.

[687] Elhamifar, E., Burden, S. A., and Sastry, S. S. (2014). Adaptive piecewise–affine inverse modeling of hybrid dynamical systems. *IFAC Proceedings Volumes*, *47*(3), 10844-10849.

[688] Elzinga, M. J., van Breugel, F., and Dickinson, M. H. (2014). Strategies for the stabilization of longitudinal forward flapping flight revealed using a dynamically-scaled robotic fly. *Bioinspiration & biomimetics*, *9*(2), 025001.

[689] Estrada, M. A., Hawkes, E. W., Christensen, D. L., and Cutkosky, M. R. (2014). Perching and crawling: design of a multimodal robot. In *IEEE International Conference on Robotics and Automation (ICRA)*.

[690] Faruque, I. A., and Humbert, J. S. (2014). Wing motion transformation to evaluate aerodynamic coupling in flapping wing flight. *Journal of theoretical biology*, *363*, 198-204.

[691] Faust, A., Palunko, I., Cruz, P., Fierro, R., and Tapia, L. (2014). Automated aerial suspended cargo delivery through reinforcement learning. *Artificial Intelligence*.

[692] Faust, A., Ruymgaart, P., Salman, M., Fierro, R., and Tapia, L. (2014). Continuous action reinforcement learning for control-affine systems with unknown dynamics. *IEEE/CAA Journal of Automatica Sinica*, *1*(3), 323-336.

[693] Fick, D., Kim, G., Wang, A., Blaauw, D., and Sylvester, D. (2014, September). Mixed-signal stochastic computation demonstrated in an image sensor with integrated 2D edge detection and noise filtering. In *Proceedings of the IEEE 2014 Custom Integrated Circuits Conference* (pp. 1-4). IEEE.

[694] Figueroa, R., Faust, A., Cruz, P., Tapia, L., and Fierro, R. (2014, June). Reinforcement learning for balancing a flying inverted pendulum. In *Intelligent Control and Automation (WCICA), 2014 11th World Congress on* (pp. 1787-1793). IEEE.

[695] Forrest, S.R., and Lee, K. (2013). Making high efficiency, flexible, semiconductor solar cells bonded to plastic substrates. In *Materials Research Society 2013 Fall Meeting*, Boston, MA. December 2-6, 2013 (Invited).

[696] Gao, H., and Wei, M. (2013, November). Model order reduction for the coupled system of flow and moving structure. In *APS Division of Fluid Dynamics Meeting Abstracts* (Vol. 1, p. 5004).

[697] Gao, H., and Wei, M. (2014). Global model reduction for flows with moving boundary. In *52nd Aerospace Sciences Meeting* (p. 0222).

[698] Ghosh, D., and Baeder, J. D. (2014). Weighted non-linear compact schemes for the direct numerical simulation of compressible, turbulent flows. *Journal of Scientific Computing*, *61*(1), 61-89.

[699] Ghosh, D., Medida, S., and Baeder, J. D. (2014). Application of compact-reconstruction weighted essentially nonoscillatory schemes to compressible aerodynamic flows. *AIAA Journal*, *52*(9), 1858-1870.

[700] Gil, S., Kumar, S., Katabi, D., and Rus, D. (2015). Adaptive communication in multi-robot systems using directionality of signal strength. *The International Journal of Robotics Research*, *34*(7), 946-968.

[701] Gokhale, V. J., and Rais-Zadeh, M. (2014). Phonon-electron interactions in piezoelectric semiconductor bulk acoustic wave resonators. *Scientific reports*, *4*.

[702] Gokhale, V. J., and Rais-Zadeh, M. (2014). Uncooled infrared detectors using gallium nitride on silicon micromechanical resonators. *Journal of Microelectromechanical Systems*, *23*(4), 803-810.

[703] Goldman, D. I. (2014). Colloquium: biophysical principles of undulatory self-propulsion in granular media. *Reviews of Modern Physics*, *86*(3), 943.

[704] Goldman, D. I. (2014). Swimming through sand: biological physics of animal movement in a complex medium. *Reviews of Modern Physics*.

[705] Gong, C., Travers, M., Astley, H., Li, L., Mendelson, J., Hu, D., and Choset, H. (2014). Conditioned basis array factorization: An approach to gait pattern extraction. *Proceedings of Robotics: Science and Systems X,* Berkeley, CA, 12-16.

[706] Gravish, N., and Goldman, D. I. (2014). Effect of volume fraction on granular avalanche dynamics. *Physical Review E*, *90*(3), 032202.

[707] Gravish, N., Gold, G., and Zangwill, A., Goodisman, M. A., and Goldman, D. I. (2014, March). Congestion and communication in confined ant traffic. *APS Meeting Abstracts*, *1*, 10003.

[708] Gravish, N., Monaenkova, D., Goodisman, M. A., and Goldman, D. I. (2013). Climbing, falling, and jamming during ant locomotion in confined environments. *Proceedings of the National Academy of Sciences*, *110*(24), 9746-9751.

[709] Gravish, N., Umbanhowar, P. B., and Goldman, D. I. (2014). Force and flow at the onset of drag in plowed granular media. *Physical Review E*, *89*(4), 042202.

[710] Grocholsky, B., and Michael, N. (2013, December). Exploiting mobility heterogeneity in micro-aerial vehicle deployments for environment exploration. In *2013 IEEE Globecom Workshops (GC Wkshps)* (pp. 1421-1425). IEEE.

[711] Haldane, D. W., and Fearing, R. S. (2014, May). Roll oscillation modulated turning in dynamic millirobots. In *2014 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 4569-4575). IEEE.

[712] Harrington, A. M., and Kroninger, C. M. (2014, June). Endurance bounds of aerial systems. In *SPIE Defense+ Security* (pp. 90831R-90831R). International Society for Optics and Photonics.

[713] Hawkes, E. W., Christensen, D. L., Eason, E. V., Estrada, M. A., Heverly, M., Hilgemann, E., and Cutkosky, M. R. (2013, November). Dynamic surface grasping with directional adhesion. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 5487-5493). IEEE.

[714] Hickson, S., Birchfield, S., Essa, I., and Christensen, H. (2014). Efficient hierarchical graph-based segmentation of RGBD videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 344-351).

[715] Hosoi, A. E., and Goldman, D. I. (2015). Beneath our feet: strategies for locomotion in granular media. *Annual Review of Fluid Mechanics*, *47*, 431-453.

[716] Hrishikeshavan, V., Benedict, M., and Chopra, I. (2014). Identification of flight dynamics of a cylcocopter micro air vehicle in hover. *Journal of Aircraft*, *52*(1), 116-129.

[717] Hunt, N. H., Moon, H. S., Hammond, A., Burnett, N. P., Pritchard-Berman, M., and Full, R. J. (2014, January). Stability of cockroaches running rapidly on rigid rods. *Integrative and Comparative Biology*, *54*, E97-E97. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[718] Indelman, V., Carlone, L., and Dellaert, F. (2014, May). Planning under uncertainty in the continuous domain: a generalized belief space approach. In *2014 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 6763-6770). IEEE.

[719] Indelman, V., Carlone, L., and Dellaert, F. (2015). Planning in the continuous domain: A generalized belief space approach for autonomous navigation in unknown environments. *The International Journal of Robotics Research*, *34*(7), 849-882.

[720] Indelman, V., Nelson, E., Michael, N., and Dellaert, F. (2014, May). Multi-robot pose graph localization and data association from unknown initial relative poses via expectation maximization. In *2014 IEEE*

International Conference on Robotics and Automation (ICRA) (pp. 593-600). IEEE.

[721] Jam, A., Moallem, M., East, J., and Sarabandi, K. (2014). A non-contact waveguide probe for on-wafer-parameter measurements for submillimeter-wave to terahertz band. *IEEE Transactions on Terahertz Science and Technology*, *4*(4), 515-522.

[722] Jam, A., Sarabandi, K., and Vahidpour, M. (2014, July). A novel frequency beam-steering antenna array at Y-band. In *2014 IEEE Antennas and Propagation Society International Symposium (APSURSI)* (pp. 1371-1372). IEEE.

[723] Jam, A., Vahidpour, M., and Sarabandi, K. (2014). A Y-band micromachined slot-coupled cavity-backed patch antenna array. In *IEEE International Symposium on Antennas and Propagation*, Memphis, TN.

[724] Jam, A., Vahidpour, M., East, J., and Sarabandi, K. (2014, August). Microfabricaion and measurement of a sub-millimeterwave beam-scanning antenna array at Y-band. In *General Assembly and Scientific Symposium (URSI GASS), 2014 XXXIth URSI* (pp. 1-2). IEEE.

[725] Jarugumilli, T., Benedict, M., and Chopra, I. (2014). Wind tunnel studies on a micro air vehicle–scale cycloidal rotor. *Journal of the American Helicopter Society*, *59*(2), 1-10.

[726] Jayaram, K., Goldman, D. I., and Full, R. J. (2014). Mechanics of confined space frictional legged crawling. *World Congress of Biomechanics*, Boston, MA.

[727] Jayaram, K., Goldman, D., and Full, R. J. (2014, January). Effect of friction on cockroaches running in confined spaces. *Integrative and Comparative Biology*, *54*, E101-E101. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[728] Jiang, H., Pope, M. T., Hawkes, E. W., Christensen, D. L., Estrada, M. A., Parlier, A., and Cutkosky, M. R. (2014, May). Modeling the dynamics of perching with opposed-grip mechanisms. In *2014 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 3102-3108). IEEE.

[729] Karydis, K., Liu, Y., Poulakakis, I., and Tanner, H. G. (2015). A template candidate for miniature legged robots in quasi-static motion. *Autonomous Robots*, *38*(2), 193-209.

[730] Karydis, K., Zarrouk, D., Poulakakis, I., Fearing, R. S., and Tanner, H. G. (2014, September). Planning with the STAR (s). In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 3033-3038). IEEE.

[731] Kessens, C. C., Lennon, C. T., and Collins, J. (2014, September). A metric for self-rightability and understanding its relationship to simple

morphologies. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 3699-3704). IEEE.

[732] Kohut, N. J., Zarrouk, D., Peterson, K. C., and Fearing, R. S. (2013, November). Aerodynamic steering of a 10 cm high-speed running robot. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 5593-5599). IEEE.

[733] Kolluru Venkata, S. & Jones, A. R. (2012). Leading edge vortex development on a pitching and rotating wing. AIAA Paper 2012-0050 in the *50th AIAA Aerospace Sciences Meeting and Exhibit*, Nashville, TN, Jan. 9-12, 2012.

[734] Kroninger, C., Harrington, A., and Munson, M. (2014). A parametric study of maximum flight efficiency for micro aerial vehicles. In *Science and Technology Forum and Exposition, AIAA*.

[735] Kumar, S., Gil, S., Katabi, D., and Rus, D. (2014, September). Accurate indoor localization with zero start-up cost. In *Proceedings of the 20th annual international conference on Mobile computing and networking* (pp. 483-494). ACM.

[736] Lenert, A., Bierman, D. M., Nam, Y., Chan, W. R., Celanovic, I., Soljacic, M., and Wang, E. N. (2014.) Nanophotonic solar thermophotovoltaics for high-performance solar power generation. *Nature Nanotechnologies*, *9*, 126-130.

[737] Li, C., Fearing, R., and Full, R. (2014, March). Insects traversing grass-like vertical compliant beams. In *APS Meeting Abstracts* (Vol. 1, p. 10012).

[738] Li, C., Pullin, A., Haldane, D., Fearing, R., and Full, R. (2014, November). Shape-assisted body reorientation enhances trafficability through cluttered terrain. In *APS Division of Fluid Dynamics Meeting Abstracts*.

[739] Libby, T., and Full, R. J. (2014). Model-based tools for decoupling the effects of internal and external forces on dynamic maneuvers. *World Congress of Biomechanics*, Boston, MA.

[740] Libby, T., Hwang, M., Koh, M., Xie, B., and Full, R. J. (2014, January). Ground reaction forces during rapid escape turns in lizards. *Integrative and Comparative Biology*, *54*, E123-E123. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[741] Liu, L., and Michael, N. (2014, May). Energy-aware aerial vehicle deployment via bipartite graph matching. In *Unmanned Aircraft Systems (ICUAS), 2014 International Conference on* (pp. 189-194). IEEE.

[742] Liu, L., Michael, N., and Shell, D. (2014, July). Fully decentralized task swaps with optimized local searching. In *Proceedings of Robotics: Science and Systems*.

[743] Lu, W., Zhang, G., Ferrari, S., Anderson, M., and Fierro, R. (2014). A particle-filter information potential method for tracking and monitoring maneuvering targets using a mobile sensor agent. *The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*, *11*(1), 47-58.

[744] Mancini, P. M., Jones, A. R., Granlund, K., and Ol, M. (2014, January). Parameter studies on translating rigid and flexible wings. In *52nd Aerospace Sciences Meeting, AIAA* (Vol. 73).

[745] Mancini, P., Jones, A., Granlund, K., and Ol, M. (2013, November). Force coefficients on surging rigid and flexible wings. In *APS Division of Fluid Dynamics Meeting Abstracts* (Vol. 1, p. 2009).

[746] Marvi, H., Gong, C., Gravish, N., Astley, H., Travers, M., Hatton, R. L., and Goldman, D. I. (2014). Sidewinding with minimal slip: Snake and robot ascent of sandy slopes. *Science*, *346*(6206), 224-229.

[747] Matthies, L., Brockers, R., Kuwata, Y., and Weiss, S. (2014, May). Stereo vision-based obstacle avoidance for micro air vehicles using disparity space. In *2014 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 3242-3249). IEEE.

[748] Mayo, D., Lankford, J., Benedict, M., and Chopra, I. (2014). Experimental and computational aerodynamic investigation of avian-based rigid flapping wings for MAV applications. Conference proceedings of the *Fifth Decennial AHS Aeromechanics Specialists' Conference*, San Francisco, CA, January 2014.

[749] Mayo, D., Lankford, J., Benedict, M., andChopra, I. (2014). Coupled CFD/CSD-based aeroelastic analysis with flowfield and airloads measurements of avian-based flexible flapping wings for MAV applications. Conference proceedings of the *70th Annual Forum and Technology Display of the American Helicopter Society*, Palais des congrès de Montréal, Montréal, Québec, May 2014.

[750] Moallem, M., and Sarabandi, K. (2014). Polarimetric study of MMW imaging radars for indoor navigation and mapping. *IEEE Transactions on Antennas and Propagation*, *62*(1), 500-504.

[751] Moallem, M., and Sarabandi, K. (2014, July). A novel non-contact s-parameter measurement method for sub-MMW multi-port on-wafer devices. In *Radio Science Meeting (Joint with AP-S Symposium), 2014 USNC-URSI* (pp. 2-2). IEEE.

[752] Moallem, M., and Sarabandi, K. (2014, July). Broadband micromachined J-band directional couplers. In *Radio Science Meeting (Joint with AP-S Symposium), 2014 USNC-URSI* (pp. 213-213). IEEE.

[753] Moallem, M., East, J., and Sarabandi, K. (2014). A broadband, micromachined rectangular waveguide to cavity-backed coplanar waveguide transition using impedance-taper technique. *IEEE Transactions on Terahertz Science and Technology*, *4*(1), 49-55.

[754] Moldovan, T. M., Jordan, M. I., and Abbeel, P. (2013). Dirichlet process reinforcement learning. In *Reinforcement Learning and Decision Making Meeting*.

[755] Muijres, F. T., Elzinga, M. J., Melis, J. M., and Dickinson, M. H. (2014). Flies evade looming targets by executing rapid visually directed banked turns. *Science*, *344*(6180), 172-177.

[756] Mukherjee, R., and Kim, I. (2013, August). Massively parallel discrete element modeling of legged mobility on granular terrain. In *ASME 2013 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (pp. V07AT10A068-V07AT10A068). American Society of Mechanical Engineers.

[757] Mulgaonkar, Y., Whitzer, M., Morgan, B., Kroninger, C. M., Harrington, A. M., and Kumar, V. (2014, June). Power and weight considerations in small, agile quadrotors. In *SPIE Defense+ Security* (pp. 90831Q-90831Q). International Society for Optics and Photonics.

[758] Nam, Y., Yeng, Y. X., Lenert, A., Bermel, P., Celanovic, I., Soljačić, M., and Wang, E. N. (2014). Solar thermophotovoltaic energy conversion systems with two-dimensional tantalum photonic crystal absorbers and emitters. *Solar Energy Materials and Solar Cells*, *122*, 287-296.

[759] Nelson, E., Indelman, V., Michael, N., and Dellaert, F. (2016). An experimental study of robust distributed multi-robot data association from arbitrary poses. In *Experimental Robotics* (pp. 323-337). Springer International Publishing.

[760] Nieto-Granda, C., Choudhary, S., Rogers, J. G., Twigg, J., Murali, V., and Christensen, H. I. (2014, June). Object guided autonomous exploration for mobile robots in indoor environments. In *SPIE Defense+ Security* (pp. 90840M-90840M). International Society for Optics and Photonics.

[761] Nieto-Granda, C., Rogers III, J. G., and Christensen, H. I. (2013, May). Multi-robot exploration strategies for tactical tasks in urban environments. In *SPIE Defense, Security, and Sensing* (pp. 87410B-87410B). International Society for Optics and Photonics.

[762] Nieto-Granda, C., Rogers, J. G., and Christensen, H. I. (2014). Coordination strategies for multi-robot exploration and mapping. *The International Journal of Robotics Research*, 0278364913515309.

[763] Nieto-Granda, C., Rogers, J. G., and Christensen, H. I. (2014). Coordination strategies for multi-robot exploration and mapping. *The International Journal of Robotics Research*, 0278364913515309.

[764] Özaslan, T., Shen, S., Mulgaonkar, Y., Michael, N., and Kumar, V. (2015). Inspection of penstocks and featureless tunnel-like environments using micro UAVs. In *Field and Service Robotics* (pp. 123-136). Springer International Publishing.

[765] Pahlajani, C. D., Sun, J., Poulakakis, I., and Tanner, H. G. (2014). Error probability bounds for nuclear detection: Improving accuracy through controlled mobility. *Automatica*, *50*(10), 2470-2481.

[766] Pahlajani, C. D., Sun, J., Poulakakis, I., and Tanner, H. G. (2015). Performance bounds for mismatched decision schemes with Poisson process observations. *Systems & Control Letters*, *75*, 69-76.

[767] Palunko, P. Cruz, and R. Fierro. Load Transportation Using Aerial Robots: Safe and Efficient Load Manipulation, Briefs in Applied Science and Technology, Springer 2014 (To appear February 2015). ISBN-13: 978-3319032269, ISBN-10: 3319032267.

[768] Park, S., Cho, J., Lee, K., and Yoon, E. (2014, February). 243.3 pJ/pixel bio-inspired time-stamp-based 2D optic flow sensor for artificial compound eyes. In *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)* (pp. 126-127). IEEE.

[769] Purohit, A., Zhang, P., Sadler, B. M., and Carpin, S. (2014, May). Deployment of swarms of micro-aerial vehicles: from theory to practice. In *2014 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 5408-5413). IEEE.

[770] Rahimi, N., Aragon, A. A., Romero, O. S., Shima, D. M., Rotter, T. J., Mukherjee, S. D., and Lester, L. F. (2013). Electrical and microstructure analysis of nickel-based low-resistance ohmic contacts to n-GaSb. *APL Materials*, *1*(6), 062105.

[771] Rahimi, N., Aragon, A. A., Romero, O. S., Shima, D. M., Rotter, T. J., Mukherjee, S. D., Lester, L. F. (2014). Low resistance palladium/molybdenum based ohmic contacts to n-GaSb grown on GaAs. *Journal of Vacuum Science & Technology B*, *32*(4), 04E108.

[772] Rahimi, N., Aragon, A. A., Shima, D. M., Hains, C., Busani, T., Lavrova, O., and Lester, L. F. (2014). Characterization of surface defects on Be-implanted GaSb. *Journal of Vacuum Science & Technology B*, *32*(4), 04E109.

[773] Rahimi, N., Behzadirad, M., Renteria, E. J., Shima, D. M., Muniz, A. J., Busani, T., and Lester, L. F. (2014, March). Beryllium implant activation

and damage recovery study in n-type GaSb. In *SPIE OPTO* (pp. 89811Q-89811Q). International Society for Optics and Photonics.

[774] Rais-Zadeh, M., Gokhale, V. J., Ansari, A., Faucher, M., Théron, D., Cordier, Y., and Buchaillot, L. (2014). Gallium nitride as an electromechanical material. *Journal of Microelectromechanical Systems*, *23*(6), 1252-1271.

[775] Reyes, L. A. V., and Tanner, H. G. (2015). Flocking, formation control, and path following for a group of mobile robots. *IEEE Transactions on Control Systems Technology*, *23*(4), 1268-1282.

[776] Rinnerbauer, V., Lenert, A., Bierman, D. M., Yeng, Y. X., Chan, W. R., Geil, R. D., and Celanovic, I. (2014). Metallic photonic crystal absorber-emitter for efficient spectral control in high-temperature solar thermophotovoltaics. *Advanced Energy Materials*, *4*(12).

[777] Rinnerbauer, V., Shen, Y., Joannopoulos, J. D., Soljačić, M., Schäffler, F., and Celanovic, I. (2014). Superlattice photonic crystal as broadband solar absorber for high temperature operation. *Optics express*, *22*(107), A1895-A1906.

[778] Romero, O. S., Aragon, A. A., Rahimi, N., Shima, D., Addamane, S., Mukherjee, S. D., Lavrova, O., Lester, L. F., and Balakrishnan, G. (2014). Surface defect based failure mechanism in large area GaSb diodes grown using molecular beam epitaxy. *Journal of Electronic Materials*, *43*(4), 926-930.

[779] Sarabandi, K., Oh, J., Pierce, L., Shivakumar, K., and Lingaiah, S. (2014). Lightweight, conformal antennas for robotic flapping flyers. *IEEE Antennas and Propagation Magazine*, *56*(6), 29-40.

[780] Saska, M., Chudoba, J., Přeučil, L., Thomas, J., Loianno, G., Třešňák, A., and Kumar, V. (2014, May). Autonomous deployment of swarms of micro-aerial vehicles in cooperative surveillance. In *Unmanned Aircraft Systems (ICUAS), 2014 International Conference on* (pp. 584-595). IEEE.

[781] Sharpe, S. S., Koehler, S. A., Kuckuk, R. M., Serrano, M., Vela, P. A., Mendelson, J., and Goldman, D. I. (2015). Locomotor benefits of being a slender and slick sand swimmer. *Journal of Experimental Biology*, *218*(3), 440-450.

[782] Sheikhsofla, M., and Sarabandi, K. (2013, July). Indoor wave propagation simulations at HF using rayleigh-gans approximation. In *Radio Science Meeting (Joint with AP-S Symposium), 2013 USNC-URSI* (pp. 202-202). IEEE.

[783] Shen, S., Mulgaonkar, Y., Michael, N., and Kumar, V. (2016). Initialization-free monocular visual-inertial state estimation with

application to autonomous MAVs. In *Experimental Robotics* (pp. 211-227). Springer International Publishing.

[784] Shen, Y., Ye, D., Celanovic, I., Johnson, S. G., Joannopoulos, J. D., and Soljačić, M. (2014). Optical broadband angular selectivity. *Science*, *343*(6178), 1499-1501.

[785] Shen, Y., Ye, D., Wang, L., Celanovic, I., Ran, L., Joannopoulos, J. D., and Soljačić, M. (2014). Metamaterial broadband angular selectivity. *Physical Review B*, *90*(12), 125422.

[786] Shrestha, E., Hrishikeshavan, V., Benedict, M., Yeo, D., and Chopra, I. (2014). Development of control strategies and flight testing of a twin-cyclocopter in forward flight. In *AHS 70th Annual Forum*.

[787] Springthorpe, D., Ng, P., and Full, R. J. (2014, January). Effect of moisture content on burrowing performance of ghost crabs. *Integrative and Comparative Biology*, *54*, E198-E198. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[788] Sreenath, K., Lee, T., and Kumar, V. (2013, December). Geometric control and differential flatness of a quadrotor UAV with a cable-suspended load. In *52nd IEEE Conference on Decision and Control* (pp. 2269-2274). IEEE.

[789] Sreenath, K., Michael, N., and Kumar, V. (2013, May). Trajectory generation and control of a quadrotor with a cable-suspended load-A differentially-flat hybrid system. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on* (pp. 4888-4895). IEEE.

[790] Stelmakh, V., Rinnerbauer, V., Chan, W. R., Senkevich, J. J., Joannopoulos, J. D., Soljačić, M., and Celanovic, I. (2014, June). Performance of tantalum-tungsten alloy selective emitters in thermophotovoltaic systems. In *SPIE Sensing Technology+ Applications* (pp. 911504-911504). International Society for Optics and Photonics.

[791] Stephan, J., Fink, J., Charrow, B., Ribeiro, A., and Kumar, V. (2014, September). Robust routing and multi-confirmation transmission protocol for connectivity management of mobile robotic teams. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 3753-3760). IEEE.

[792] Tanner, H. G. (2014). Relaxed stability conditions for switched systems with dwell time. *Asian Journal of Control*, *16*(2), 313-320.

[793] Thomas, J., Loianno, G., Polin, J., Sreenath, K., and Kumar, V. (2014). Toward autonomous avian-inspired grasping for micro aerial vehicles. *Bioinspiration & biomimetics*, *9*(2), 025010.

[794] Thomas, J., Loianno, G., Sreenath, K., and Kumar, V. (2014, May). Toward image based visual servoing for aerial grasping and perching. In

2014 IEEE International Conference on Robotics and Automation (ICRA) (pp. 2113-2118). IEEE.

[795] Tran, J., Sirohi, J., Gao, H., and Wei, M. (2015). Reduced-order modeling of loads and deformation of a flexible flapping wing. *Target*, *20*, 40.

[796] Trevor, A. J., Rogers, J. G., and Christensen, H. I. (2014, May). Omnimapper: A modular multimodal mapping framework. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on* (pp. 1983-1990). IEEE.

[797] Trevor, A. J., Rogers, J. G., and Christensen, H. I. (2014, May). Omnimapper: A modular multimodal mapping framework. In *2014 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 1983-1990). IEEE.

[798] Turpin, M., Michael, N., and Kumar, V. (2014). Capt: concurrent assignment and planning of trajectories for multiple robots. *The International Journal of Robotics Research*, *33*(1), 98-112.

[799] Turpin, M., Michael, N., and Kumar, V. (2015). An approximation algorithm for time optimal multi-robot routing. In *Algorithmic Foundations of Robotics XI* (pp. 627-640). Springer International Publishing.

[800] Valbuena, L., Cruz, P., Figueroa, R., Sorrentino, F., and Fierro, R. (2014, September). Stable formation of groups of robots via synchronization. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 376-381). IEEE.

[801] Venkata, S. K., and Jones, A. R. (2013). Leading-edge vortex structure over multiple revolutions of a rotating wing. *Journal of Aircraft*, *50*(4), 1312-1316.

[802] Yektakhah, B., and Sarabandi, K. (2014, July). All-directions through the wall radar imaging using a synthetic array formed by a finite number of Moving Robots. In *2014 IEEE Antennas and Propagation Society International Symposium (APSURSI)* (pp. 1111-1112). IEEE.

[803] Yeng, Y. X., Chou, J. B., Rinnerbauer, V., Shen, Y., Kim, S. G., Joannopoulos, J. D., and Čelanović, I. (2014). Global optimization of omnidirectional wavelength selective emitters/absorbers based on dielectric-filled anti-reflection coated two-dimensional metallic photonic crystals. *Optics express*, *22*(18), 21711-21718.

[804] Zhang, T., and Goldman, D. I. (2014). The effectiveness of resistive force theory in granular locomotiona). *Physics of Fluids (1994-present)*, *26*(10), 101308.

[805] Aguilar, J., and Goldman, D. I. (2015). Robophysical study of jumping dynamics on granular media. *Nature Physics*.

[806] Ansari, A., and Rais-Zadeh, M. (2014). A temperature-compensated micromechanical gallium nitride resonator. *IEEE Electron Device Letters*, *35*(11), 1127-1129.

[807] Ansari, A., and Rais-Zadeh, M. (2016, January). Frequency-tunable current-assisted AlGaN/GaN acoustic resonators. In *2016 IEEE 29th International Conference on Micro Electro Mechanical Systems (MEMS)* (pp. 123-126). IEEE.

[808] Astley, H. C., Gong, C., Dai, J., Travers, M., Serrano, M. M., Vela, P. A., and Goldman, D. I. (2015). Modulation of orthogonal body waves enables high maneuverability in sidewinding locomotion. *Proceedings of the National Academy of Sciences*, *112*(19), 6200-6205.

[809] Badrya, C., and Baeder, J. D. (2015). Insect kinematics in trimmed flight at low Reynolds numbers using CFD. In *53rd AIAA Aerospace Sciences Meeting* (p. 0045).

[810] Bry, A., Richter, C., Bachrach, A., and Roy, N. (2015). Aggressive flight of fixed-wing and quadrotor aircraft in dense indoor environments. *The International Journal of Robotics Research*, *34*(7), 969-1002.

[811] Cai, X., and Sarabandi, K. (2015, July). A compact broadband horizontally polarized omnidirectional antenna using folded dipoles. In *Radio Science Meeting (Joint with AP-S Symposium), 2015 USNC-URSI* (pp. 9-9). IEEE.

[812] Cai, X., and Sarabandi, K. (2016). A compact broadband horizontally polarized omnidirectional antenna using planar folded dipole elements. *IEEE Transactions on Antennas and Propagation*, *64*(2), 414-422.

[813] Chen, X., Purohit, A., Dominguez, C. R., Carpin, S., and Zhang, P. (2015, November). DrunkWalk: collaborative and adaptive planning for navigation of micro-aerial sensor swarms. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems* (pp. 295-308). ACM.

[814] Choi, J., Dagefu, F. T., Sadler, B. M., and Sarabandi, K. (2016). Electrically small folded dipole antenna for HF and low-VHF bands. *IEEE Antennas and Wireless Propagation Letters*, *15*, 718-721.

[815] Choi, J., Park, S., Cho, J., and Yoon, E. (2015). An energy/illumination-adaptive CMOS image sensor with reconfigurable modes of operations. *IEEE Journal of Solid-State Circuits*, *50*(6), 1438-1450.

[816] Choi, J., Sabet, A., Sabet, K., and Sarabandi, K. (2015, July). Near-field characterizations of a miniaturized low-VHF omni-directional antenna

using an electro-optical system. In *Radio Science Meeting (Joint with AP-S Symposium), 2015 USNC-URSI* (pp. 69-69). IEEE.

[817] Choudhary, S., Carlone, L., Christensen, H. I., and Dellaert, F. (2015, September). Exactly sparse memory efficient SLAM using the multi-block alternating direction method of multipliers. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on* (pp. 1349-1356). IEEE.

[818] Choudhary, S., Indelman, V., Christensen, H. I., and Dellaert, F. (2015, May). Information-based reduced landmark SLAM. In *2015 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 4620-4627). IEEE.

[819] Choudhary, S., Trevor, A. J., Christensen, H. I., and Dellaert, F. (2014, September). SLAM with object discovery, modeling and mapping. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 1018-1025). IEEE.

[820] Chow, Y. L., Pavone, M., Sadler, B. M., and Carpin, S. (2015). Trading safety versus performance: Rapid deployment of robotic swarms with robust performance constraints. *Journal of Dynamic Systems, Measurement, and Control*, *137*(3), 031005.

[821] Conroy, J., Gremillion, G., Mathis, A., and Nothwang, W. (2015, May). An analysis of small target feature detection limits using optic flow. In *SPIE Defense+ Security* (pp. 94671Q-94671Q). International Society for Optics and Photonics.

[822] Cruz, P. J., and Fierro, R. (2015, December). Building coalitions of heterogeneous agents using weighted bipartite graphs. In *2015 54th IEEE Conference on Decision and Control (CDC)* (pp. 2822-2828). IEEE.

[823] Cruz, P. J., and Fierro, R. (2015, June). Toward optical wireless communications between micro unmanned aerial and ground systems. In *Unmanned Aircraft Systems (ICUAS), 2015 International Conference on* (pp. 669-676). IEEE.

[824] Cruz, P. J., Fierro, R., and Abdallah, C. T. (2016). Cooperative Learning for Robust Connectivity in Multirobot Heterogeneous Networks. *Control of Complex Systems: Theory and Applications*, 451.

[825] Cruz, P. J., Oishi, M., and Fierro, R. (2015, July). Lift of a cable-suspended load by a quadrotor: A hybrid system approach. In *2015 American Control Conference (ACC)* (pp. 1887-1892). IEEE.

[826] Cruz, P., and Fierro, R. (2014, October). Autonomous lift of a cable-suspended load by an unmanned aerial robot. In *2014 IEEE Conference on Control Applications (CCA)* (pp. 802-807). IEEE.

[827] Dagefu, F. T., Choi, J., Sheikhsofla, M., Sadler, B. M., and Sarabandi, K. (2015). Performance assessment of lower VHF band for short-range communication and geolocation applications. *Radio Science*, *50*(5), 443-452.

[828] Dagefu, F. T., Verma, G., Rao, C. R., Paul, L. Y., Fink, J. R., Sadler, B. M., and Sarabandi, K. (2015). Short-range Low-VHF channel characterization in cluttered environments. *IEEE Transactions on Antennas and Propagation*, *63*(6), 2719-2727.

[829] Dagefu, F. T., Verma, G., Rao, C. R., Paul, L. Y., Sadler, B. M., and Sarabandi, K. (2015, March). Measurement and characterization of the short-range low-VHF channel. In *2015 IEEE Wireless Communications and Networking Conference (WCNC)* (pp. 177-182). IEEE.

[830] Dayanik, M. B., Collins, N., and Flynn, M. P. (2015, September). A 28.5–33.5 GHz fractional-N PLL using a 3rd order noise shaping time-to-digital converter with 176fs resolution. In *European Solid-State Circuits Conference (ESSCIRC), ESSCIRC 2015-41st* (pp. 376-379). IEEE.

[831] Desaraju, V. R., Michael, N., Humenberger, M., Brockers, R., Weiss, S., Nash, J., and Matthies, L. (2015). Vision-based landing site evaluation and informed optimal trajectory generation toward autonomous rooftop landing. *Autonomous Robots*, *39*(3), 445-463.

[832] Dimble, K., Escobar-Alvarez, H., Ranganathan, B., Conroy, J., and Humbert, J. S. (2015, January 20-22). 3D depth estimation for helicopter landing site visualization in environments with degraded visibility. In *Proceedings of the American Helicopter Society Specialist's Meeting on Unmanned Rotorcraft Systems*.

[833] Estrada, M. A., Hawkes, E. W., Christensen, D. L., and Cutkosky, M. R. (2014, May). Perching and vertical climbing: Design of a multimodal robot. In *2014 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 4215-4221). IEEE.

[834] Faust, A., Palunko, I., Cruz, P., Fierro, R., and Tapia, L. (2014). Automated aerial suspended cargo delivery through reinforcement learning. *Artificial Intelligence*.

[835] Fu, C., Arguello, H., Sadler, B. M., and Arce, G. R. (2015). Compressive spectral polarization imaging by a pixelized polarizer and colored patterned detector. *JOSA A*, *32*(11), 2178-2188.

[836] Gao, H., and Wei, M. (2014, November). Global model reduction for the aerodynamics of coupled fluid-structure systems. In *APS Meeting Abstracts*.

[837] Gao, H., and Wei, M. (2016). Domain decomposition in POD-Galerkin projection for flows with moving boundary. In *54th AIAA Aerospace Sciences Meeting* (p. 1102).

[838] Gil, S., Kumar, S., Katabi, D., and Rus, D. (2015). Adaptive communication in multi-robot systems using directionality of signal strength. *The International Journal of Robotics Research*, *34*(7), 946-968.

[839] Gil, S., Kumar, S., Mazumder, M., Katabi, D., and Rus, D. (2015). Guaranteeing spoof-resilient multi-robot networks.

[840] Gokhale, V. J., Figueroa, C., Tsai, J. M. L., and Rais-Zadeh, M. (2015, January). Low-noise AlN-on-Si resonant infrared detectors using a commercial foundry MEMS fabrication process. In *2015 28th IEEE International Conference on Micro Electro Mechanical Systems (MEMS)* (pp. 73-76). IEEE.

[841] Gokhale, V. J., Myers, P. D., and Rais-Zadeh, M. (2014, November). Subwavelength plasmonic absorbers for spectrally selective resonant infrared detectors. In *IEEE SENSORS 2014 Proceedings* (pp. 982-985). IEEE.

[842] Gong, C., Travers, M. J., Astley, H. C., Li, L., Mendelson, J. R., Goldman, D. I., and Choset, H. (2016). Kinematic gait synthesis for snake robots. *The International Journal of Robotics Research*, *35*(1-3), 100-113.

[843] Gremillion, G. M., Castano, L. M., and Humbert, J. S. (2015). Disturbance rejection with distributed acceleration and strain sensing. In *AIAA SciTech Conference*.

[844] Gremillion, G., and Humbert, J. S. (2015, May). Bio-inspired sensing and control for disturbance rejection and stabilization. In *SPIE Defense+ Security* (pp. 94671O-94671O). International Society for Optics and Photonics.

[845] Haldane, D. W., and Fearing, R. S. (2015, May). Running beyond the bio-inspired regime. In *2015 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 4539-4546). IEEE.

[846] Haldane, D. W., Casarez, C. S., Karras, J. T., Lee, J., Li, C., Pullin, A. O., and Fearing, R. S. (2015). Integrated manufacture of exoskeletons and sensing structures for folded millirobots. *Journal of Mechanisms and Robotics*, *7*(2), 021011.

[847] Hawkes, E. W., Christensen, D. L., and Cutkosky, M. R. (2015, May). Vertical dry adhesive climbing with a 100× bodyweight payload. In *2015 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 3762-3769). IEEE.

[848] Hawkes, E. W., Christensen, D. L., Han, A. K., Jiang, H., and Cutkosky, M. R. (2015, May). Grasping without squeezing: Shear adhesion gripper

with fibrillar thin film. In *2015 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 2305-2312). IEEE.

[849] Hawkes, E. W., Jiang, H., and Cutkosky, M. R. (2015). Three-dimensional dynamic surface grasping with dry adhesion. *The International Journal of Robotics Research*, 0278364915584645.

[850] Hosoi, A. E., and Goldman, D. I. (2015). Beneath our feet: strategies for locomotion in granular media. *Annual Review of Fluid Mechanics*, *47*, 431-453.

[851] Hrishikeshavan, V., and Chopra, I. (2015). Development of a lightweight inertial navigation system for micro air vehicle applications. In *Proceedings of the 6th AHS International Specialists' Meeting On Unmanned Rotorcraft Systems*, Chandler, AZ, January 20-22, 2015.

[852] Jam, A., East, J., and Sarabandi, K. (2015). An in-line chip integration and packaging method for millimeter- and sub-millimeter-wave band applications. In *IEEE Ant. And Propagation Symposium*.

[853] Jam, A., East, J., and Sarabandi, K. (2015, May). A novel method for chip integration and packaging for millimeter-wave to terahertz band applications. In *2015 IEEE MTT-S International Microwave Symposium* (pp. 1-3). IEEE.

[854] Jayaram, K., and Full, R. J. (2015, April). Body size limit predictions for mechanically mediated maneuvers. *Integrative and Comparative Biology*, *55*, E88-E88. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[855] Jayaram, K., and Full, R.J. (2015). Compressibility enables cockroach-inspired robot to crawl rapidly in confined spaces. In *International Conference on Robotics and Automation (ICRA)*, Seattle, May 26-30, 2015.

[856] Jeon, D., Dong, Q., Kim, Y., Wang, X., Chen, S., Yu, H., and Sylvester, D. (2015, June). A 23mW face recognition accelerator in 40nm CMOS with mostly-read 5T memory. In *2015 Symposium on VLSI Circuits (VLSI Circuits)* (pp. C48-C49). IEEE.

[857] Jiang, H., Hawkes, E. W., Arutyunov, V., Tims, J., Fuller, C., King, J. P., and Cutkosky, M. R. (2015, May). Scaling controllable adhesives to grapple floating objects in space. In *2015 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 2828-2835). IEEE.

[858] Jiang, H., Pope, M. T., Estrada, M. A., Edwards, B., Cuson, M., Hawkes, E. W., and Cutkosky, M. R. (2015, September). Perching failure detection and recovery with onboard sensing. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on* (pp. 1264-1270). IEEE.

[859] Jiang, H., Pope, M. T., Hawkes, E. W., Christensen, D. L., Estrada, M. A., Parlier, A., and Cutkosky, M. R. (2014, May). Modeling the dynamics of perching with opposed-grip mechanisms. In *2014 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 3102-3108). IEEE.

[860] Kardis, K., Poulakakis, I., Sun, J., and Tanner, H. G. (2015). Probabilistic model validation for systems with uncertainty. *The International Journal of Robotics Research*.

[861] Karydis, K., Liu, Y., Poulakakis, I., and Tanner, H. G. (2015). A template candidate for miniature legged robots in quasi-static motion. *Autonomous Robots*, *38*(2), 193-209.

[862] Karydis, K., Liu, Y., Poulakakis, I., and Tanner, H. G. (2015, June). Navigation of miniature legged robots using a new template. In *Control and Automation (MED), 2015 23th Mediterranean Conference on* (pp. 1112-1117). IEEE.

[863] Kingsbury, M., Zhang, T., and Goldman, D. I. (2015). The role of foot-ground interaction and slip during bipedal walking on granular media. *Dynamic Walking*.

[864] Lankford, J., Mayo, D., and Chopra, I. (2015, January). Experimental and computational investigation of an insect based flapping wing for MAV Applications. Conference proceedings of the *6th AHS International Specialists Meeting on Unmanned Rotorcraft Systems*, Scottsdale, AZ.

[865] Lankford, J., Mayo, D., and Chopra, I. (2016). Computational investigation of insect-based flapping wings for micro air vehicle applications. *International Journal of Micro Air Vehicles*, 1756829316646640.

[866] Lee, J. S., and Fearing, R. S. (2015, May). Anisotropic collapsible leg spines for increased millirobot traction. In *2015 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 4547-4553). IEEE.

[867] Levine, S., and Abbeel, P. (2014). Learning neural network policies with guided policy search under unknown dynamics. In *Advances in Neural Information Processing Systems* (pp. 1071-1079).

[868] Levine, S., Wagener, N., and Abbeel, P. (2015, May). Learning contact-rich manipulation skills with guided policy search. In *2015 IEEE international conference on robotics and automation (ICRA)* (pp. 156-163). IEEE.

[869] Li, C., Pullin, A. O., Haldane, D. W., Lam, H. K., Fearing, R. S., and Full, R. J. (2015). Terradynamically streamlined shapes in animals and robots enhance traversability through densely cluttered terrain. *Bioinspiration & biomimetics*, *10*(4), 046003.

[870] Li, C., Pullin, A., Haldane, D., Fearing, R., and Full, R. (2014, November). Shape-assisted body reorientation enhances trafficability through cluttered terrain. In *APS Division of Fluid Dynamics Meeting Abstracts*.

[871] Libby, T., Johnson, A. M., Full, R. J., and Koditschek, D. E. (2015). Comparative morphology of inertial reorientation. In *International Conference on Robotics and Automation (ICRA)*, Seattle, May 26-30, 2015.

[872] Loianno, G., Cross, G., Qu, C., Mulgaonkar, Y., Hesch, J. A., and Kumar, V. (2015). Flying smartphones: Automated flight enabled by consumer electronics. *IEEE Robotics & Automation Magazine*, *22*(2), 24-32.

[873] Loianno, G., Mulgaonkar, Y., Brunner, C., Ahuja, D., Ramanandan, A., Chari, M., and Kumar, V. (2015, September). Smartphones power flying robots. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on* (pp. 1256-1263). IEEE.

[874] Mayo, D. B., Lankford, J. L., Benedict, M., and Chopra, I. (2015). Aeroelastic analysis of avian-based flexible flapping wings for micro air vehicles. *Journal of the American Helicopter Society*, *60*(3), 1-18.

[875] Mayo, D. B., Lankford, J. L., Benedict, M., and Chopra, I. (2015). Experimental and computational analysis of rigid flapping wings for micro air vehicles. *Journal of Aircraft*, *52*(4), 1161-1178.

[876] Meier, D., Brockers, R., Matthies, L., Siegwart, R., and Weiss, S. (2015, September). Detection and characterization of moving objects with aerial vehicles using inertial-optical flow. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on* (pp. 2473-2480). IEEE.

[877] Miao, Z., Wang, Y., and Fierro, R. (2015, December). Collision-free consensus via order preservation in multi-agent networks. In *2015 54th IEEE Conference on Decision and Control (CDC)* (pp. 2999-3005). IEEE.

[878] Miao, Z., Wang, Y., and Fierro, R. (2016). Collision-free consensus in multi-agent networks: A monotone systems perspective. *Automatica*, *64*, 217-225.

[879] Monaenkova, D., Gravish, N., Rodriguez, G., Kutner, R., Goodisman, M. A., and Goldman, D. I. (2015). Behavioral and mechanical determinants of collective subsurface nest excavation. *Journal of Experimental Biology*, *218*(9), 1295-1305.

[880] Mukherjee, R., Moreland, S., Kim, I., Lele, N., Goodwin, S., Houlihan, R., and Cutkosky, M. (2014, August). Computational and experimental approach to understanding legged mobility in micro ground vehicles. In *ASME 2014 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (pp.

V006T10A005-V006T10A005). American Society of Mechanical Engineers.

[881] Nelson, E., and Michael, N. (2015, September). Information-theoretic occupancy grid compression for high-speed information-based exploration. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on* (pp. 4976-4982). IEEE.

[882] Nieto, C., Choudhary, S., Rogers III, J. G., Twigg, J., Murali, V., and Christensen, H. I. (2015). Towards contextual awareness in robot mapping: extracting semantic hierarchy from point cloud data. In *SPIE DSS*.

[883] Pahlajani, C. D., Sun, J., Poulakakis, I., and Tanner, H. G. (2014). Error probability bounds for nuclear detection: Improving accuracy through controlled mobility. *Automatica*, *50*(10), 2470-2481.

[884] Pahlajani, C. D., Sun, J., Poulakakis, I., and Tanner, H. G. (2015). Performance bounds for mismatched decision schemes with Poisson process observations. *Systems & Control Letters*, *75*, 69-76.

[885] Panagou, D., Tanner, H. G., and Kyriakopoulos, K. J. (2015). Nonholonomic control design via reference vector fields and output regulation. *ASME Journal of Dynamic Systems, Measurement and Control*, *137*(8), 081011-081011-9.

[886] Portugal, D., Pippin, C., Rocha, R. P., and Christensen, H. (2014, September). Finding optimal routes for multi-robot patrolling in generic graphs. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 363-369). IEEE.

[887] Qian, F., and Goldman, D. (2015, May). Anticipatory control using substrate manipulation enables trajectory control of legged locomotion on heterogeneous granular media. In *SPIE Defense+ Security* (pp. 94671U-94671U). International Society for Optics and Photonics.

[888] Qian, F., and Goldman, D. I. (2015). The dynamics of legged locomotion in heterogeneous terrain: universality in scattering and sensitivity to initial conditions. In *Robotics: Science and Systems*.

[889] Qian, F., Zhang, T., Korff, W., Umbanhowar, P. B., Full, R. J., and Goldman, D. I. (2015). Principles of appendage design in robots and animals determining terradynamic performance on flowable ground. *Bioinspiration & Biomimetics*, *10*(5), 056014.

[890] Rais-Zadeh, M., Gokhale, V. J., Ansari, A., Faucher, M., Théron, D., Cordier, Y., and Buchaillot, L. (2014). Gallium nitride as an electromechanical material. *Journal of Microelectromechanical Systems*, *23*(6), 1252-1271.

[891] Reyes, L. A. V., and Tanner, H. G. (2015). Flocking, formation control, and path following for a group of mobile robots. *IEEE Transactions on Control Systems Technology*, *23*(4), 1268-1282.

[892] Richter, C., Vega-Brown, W., and Roy, N. (2015). Bayesian learning for safe high-speed navigation in unknown environments. In *Proceedings of the International Symposium on Robotics Research (ISRR 2015), Sestri Levante, Italy*.

[893] Rinnerbauer, V., Shen, Y., Joannopoulos, J. D., Soljačić, M., Schäffler, F., and Celanovic, I. (2014). Superlattice photonic crystal as broadband solar absorber for high temperature operation. *Optics Express*, *22*(107), A1895-A1906.

[894] Sadeghi, M. M., Kim, H. S., Peterson, R. L. B., and Najafi, K. (2016). Electrostatic micro-hydraulic systems. *J. Microelectromechanical Systems*, *25*(3), 557-569.

[895] Shah, S. K., and Tanner, H. G. (2015, June). Dynamics-compatible potential fields using stochastic perturbations. In *Control and Automation (MED), 2015 23th Mediterranean Conference on* (pp. 278-283). IEEE.

[896] Shah, S. K., and Tanner, H. G. (2015, May). Control of stochastic unicycle-type robots. In *2015 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 389-394). IEEE.

[897] Shah, S. K., Tanner, H. G., and Pahlajani, C. D. (2015). Optimal navigation for vehicles with stochastic dynamics. *IEEE Transactions on Control Systems Technology*, *23*(5), 2003-2009.

[898] Sharpe, S. S., Koehler, S. A., Kuckuk, R. M., Serrano, M., Vela, P. A., Mendelson, J., and Goldman, D. I. (2015). Locomotor benefits of being a slender and slick sand swimmer. *Journal of Experimental Biology*, *218*(3), 440-450.

[899] Sharpe, S. S., Kuckuk, R., and Goldman, D. I. (2015). Controlled preparation of wet granular media reveals limits to lizard burial ability. *Physical biology*, *12*(4), 046009.

[900] Shen, Y., Rinnerbauer, V., Wang, I., Stelmakh, V., Joannopoulos, J. D., and Soljačić, M. (2015). Structural colors from Fano resonances. *ACS Photonics*, *2*(1), 27-32.

[901] Shiari, B., Sadeghi, M. M., Darvishian, A., and Najafi, K. (2014, November). A discrete model for an electrostatically driven micro-hydraulic actuator. In *ASME 2014 International Mechanical Engineering Congress and Exposition* (pp. V010T13A031-V010T13A031). American Society of Mechanical Engineers.

[902] Shrestha, E., Hrishikeshavan, V., Benedict, M., Yeo, D., and Chopra, I. (2016). Development of control strategies for a twin-cyclocopter in forward flight. *Journal of the American Helicopter Society*, *61*(4), 1-9.

[903] Shrestha, E., Hrishikeshavan, V., Yeo, D., Benedict, M., and Chopra, I. (2015). Flight dynamics modeling and system identification of cyclocopter in forward flight. In *Proceedings of the American Helicopter Society 71st Forum*, Virginia Beach, VA, May 5-7, 2015.

[904] Shrestha, E., Martz, V., Yeo, D., Benedict, M., and Chopra, I. (2015). Design and hover testing of a 60-gram cyclocopter. In *Proceedings of the 6th AHS International Specialists' Meeting on Unmanned Rotorcraft Systems*, Chandler, AZ, January 20-22, 2015.

[905] Shrestha, E., Yeo, D., Hrishikeshavan, V., and Chopra, I. (2015). Autonomous maneuvering of cyclocopter using closed-loop motion capture feedback controller. In *Proceedings of the 6th AHS International Specialists' Meeting on Unmanned Rotorcraft Systems*, Chandler, AZ, January 20-22, 2015.

[906] Stager, A., Karydis, K., and Tanner, H. G. (2015, May). A passively sprawling miniature legged robot. In *2015 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 3134-3139). IEEE.

[907] Stelmakh, V., Chan, W. R., Ghebrebrhan, M., Senkevich, J., Joannopoulos, J. D., Soljačić, M., and Celanović, I. (2016). Sputtered tantalum photonic crystal coatings for high-temperature energy conversion applications. *IEEE Transactions on Nanotechnology*, *15*(2), 303-309.

[908] Sun, J., and Tanner, H. G. (2015). Constrained decision-making for low-count radiation detection by mobile sensors. *Autonomous Robots*, *39*(4), 519-536.

[909] Sun, J., Tanner, H. G., and Poulakakis, I. (2015, May). Active sensor networks for nuclear detection. In *2015 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 3549-3554). IEEE.

[910] Suresh, S. A., Christensen, D. L., Hawkes, E. W., and Cutkosky, M. (2015). Surface and shape deposition manufacturing for the fabrication of a curved surface gripper. *Journal of Mechanisms and Robotics*, *7*(2), 021005.

[911] Tabandeh, M., Wei, M., and Collins, J. P. (2016). On the symmetrization in POD-Galerkin model for linearized compressible flows. In *54th AIAA Aerospace Sciences Meeting* (p. 1106).

[912] Tang, Y., and Najafi, K. (2015, December). A two-gap capacitive structure for high aspect-ratio capacitive sensor arrays. In *2015 IEEE International Electron Devices Meeting (IEDM)* (pp. 18-2). IEEE.

[913] Thomas, J., Loianno, G., Pope, M., Hawkes, E. W., Estrada, M. A., Jiang, H., and Kumar, V. (2015, August). Planning and control of aggressive

maneuvers for perching on inclined and vertical surfaces. In *ASME 2015 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (pp. V05CT08A012-V05CT08A012). American Society of Mechanical Engineers.

[914] Tokekar, P., and Kumar, V. (2015, September). Visibility-based persistent monitoring with robot teams. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on* (pp. 3387-3394). IEEE.

[915] Tran, J., Sirohi, J., Gao, H., and Wei, M. (2015). Reduced order modeling of loads and deformation of a flexible flapping wing. In *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference* (p. 0177).

[916] Tran, J., Sirohi, J., Gao, H., and Wei, M. (2015). Reduced-order modeling of loads and deformation of a flexible flapping wing. *Target*, *20*, 40.

[917] Tron, R., Thomas, J., Loianno, G., Polin, J., Kumar, V., and Daniilidis, K. (2014). Vision-based formation control of aerial vehicles. In *Robotics: Science and Systems*.

[918] Tsai, J. M., Daneman, M., Boser, B., Horsley, D., Rais-Zadeh, M., Tang, H. Y., and Assaderaghi, F. (2015, June). Versatile CMOS-MEMS integrated piezoelectric platform. In *2015 Transducers-2015 18th International Conference on Solid-State Sensors, Actuators and Microsystems (TRANSDUCERS)* (pp. 2248-2251). IEEE.

[919] Watterson, M., and Kumar, V. (2015, September). Safe receding horizon control for aggressive MAV flight with limited range sensing. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on* (pp. 3235-3240). IEEE.

[920] Weiss, S., Brockers, R., Albrektsen, S., and Matthies, L. (2015, January). Inertial optical flow for throw-and-go micro air vehicles. In *2015 IEEE Winter Conference on Applications of Computer Vision* (pp. 262-269). IEEE.

[921] Wu, X. A., Suresh, S. A., Jiang, H., Ulmen, J. V., Hawkes, E. W., Christensen, D. L., and Cutkosky, M. R. (2015, September). Tactile sensing for gecko-inspired adhesion. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on* (pp. 1501-1507). IEEE.

[922] Xu, M., Wei, M., Yang, T., and Lee, Y. S. (2016). An embedded boundary approach for the simulation of a flexible flapping wing at different density ratio. *European Journal of Mechanics-B/Fluids*, *55*, 146-156.

[923] Xu, M., Wei, M., Yang, T., and Lee, Y. S. (2016). An embedded boundary approach for the simulation of a flexible flapping wing at different density ratio. *European Journal of Mechanics-B/Fluids*, *55*, 146-156.

[924] Yektakhah, B., and Sarabandi, K. (2016). All-directions through-the-wall radar imaging using a small number of moving transceivers. *IEEE Transactions on Geoscience and Remote Sensing*, *54*(11), 6415-6428.

[925] Yektakhah, B., and Sarabandi, K. All-directions through the wall radar imaging enhancement using orthogonal polarizations and generalized pencil of function method.

[926] Yeng, Y. X., Chan, W. R., Rinnerbauer, V., Stelmakh, V., Senkevich, J. J., Joannopoulos, J. D., and Čelanović, I. (2015). Photonic crystal enhanced silicon cell based thermophotovoltaic systems. *Optics express*, *23*(3), A157-A168.

[927] Yeo, D., Shrestha, E., Paley, D. A., and Atkins, E. (2015). Experimental development of a rotorcraft UAV downwash model for real-time disturbance localization and avoidance. In *Proceedings of AIAA Atmospheric Flight Mechanics Conference*, number AIAA 2015-1685, pages 1–14, Orlando, FL, January 2015.

[928] Zarrouk, D., Haldane, D. W., and Fearing, R. S. (2015, May). Dynamic legged locomotion for palm-size robots. In *SPIE Defense+ Security* (pp. 94671S-94671S). International Society for Optics and Photonics.

FY16

[929] Aguilar, J., and Goldman, D. I. (2015). Robophysical study of jumping dynamics on granular media. *Nature Physics*.

[930] Aguilar, J., Zhang, T., Qian, F., Kingsbury, M., McInroe, B., Mazouchova, N., and Hatton, R. L. (2016). A review on locomotion robophysics: the study of movement at the intersection of robotics, soft matter and dynamical systems. *Reports on Progress in Physics*, *79*(11), 110001.

[931] Andrews, G., Shrestha, E., and Chopra, I. (2016). Design and fabrication of a meso-scale aircraft using a cycloidal rotor propulsion system. In *54th AIAA Aerospace Sciences Meeting* (p. 0141).

[932] Ansari, A., and Rais-Zadeh, M. (2016, January). Frequency-tunable current-assisted AlGaN/GaN acoustic resonators. In *Micro Electro Mechanical Systems (MEMS), 2016 IEEE 29th International Conference on* (pp. 123-126). IEEE.

[933] Badrya, C., and Baeder, J. D. (2015). Insect kinematics in trimmed flight at low Reynolds numbers using CFD. In *53rd AIAA Aerospace Sciences Meeting* (p. 0045).

[934] Badrya, C., and Baeder, J. D. (2015, November). Numerical and experimental investigation of flow structures during insect flight. In *APS Meeting Abstracts*.

[935] Badrya, C., Harrington, A., Baeder, D. J., Kroningers, M. C., and Govindarajan, G. (2016) Prediction and validation of CFD-coupled trim analysis for a flapping wing MAV. *AHS Technical Meeting on Aeromechanics Design for Vertical Lift*.

[936] Badrya, C., Sridharan, A., and Baeder, J. D. (2016). CFD-coupled trim analysis of a millimeter-scale flapping wing MAV in steady turning flight. In *34th AIAA Applied Aerodynamics Conference* (p. 3557).

[937] Benedict, M., Jarugumilli, T., and Chopra, I. (2016). Effects of asymmetric blade-pitching kinematics on forward-flight performance of a micro-air-vehicle-scale cycloidal-rotor. *Journal of Aircraft*, *53*(5), 1568-1573.

[938] Benedict, M., Mullins, J., Hrishikeshavan, V., and Chopra, I. (2016). Development of a quad cycloidal-rotor unmanned aerial vehicle. *Journal of the American Helicopter Society*, *61*(2), 1-12.

[939] Brockers, R., Fragoso, A., and Matthies, L. (2016, May). Stereo vision-based obstacle avoidance for micro air vehicles using an egocylindrical image space representation. In *SPIE Defense+ Security* (pp. 98361R-98361R). International Society for Optics and Photonics.

[940] Cai, X., and Sarabandi, K. (2016, June). Broadband omni-directional circularly polarized antenna based on vertically and horizontally polarized elements. In *Antennas and Propagation (APSURSI), 2016 IEEE International Symposium on* (pp. 1793-1794). IEEE.

[941] Carpin, S., Chow, Y. L., and Pavone, M. (2016, May). Risk aversion in finite Markov decision processes using total cost criteria and average value at risk. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on* (pp. 335-342). IEEE.

[942] Casarez, C. S., and Fearing, R. S. (2016, May). Step climbing cooperation primitives for legged robots with a reversible connection. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on* (pp. 3791-3798). IEEE.

[943] Chan, W. R., Stelmakh, V., Allmon, W. R., Waits, C. M., Soljacic, M., Joannopoulos, J. D., and Celanovic, I. (2016, November). An integrated microcombustor and photonic crystal emitter for thermophotovoltaics. *Journal of Physics: Conference Series*, *773*(1), 012108. IOP Publishing.

[944] Chan, W. R., Stelmakh, V., Soljacic, M., Joannopoulos, J. D., Celanovic, I., and Waits, C. M. (2016, January). A thermophotovoltaic system using a photonic crystal emitter. In *ASME 2016 5th International Conference on Micro/Nanoscale Heat and Mass Transfer* (pp. V001T06A005-V001T06A005). American Society of Mechanical Engineers.

[945] Chan, W. R., Stelmakh, V., Waits, C. M., Soljacic, M., Joannopoulos, J. D., and Celanovic, I. (2015). Photonic crystal enabled thermophotovoltaics for a portable microgenerator. *Journal of Physics: Conference Series*, *660*(1), 012069. IOP Publishing.

[946] Chen, X., Purohit, A., Dominguez, C. R., Carpin, S., and Zhang, P. (2015, November). Drunkwalk: collaborative and adaptive planning for navigation of micro-aerial sensor swarms. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems* (pp. 295-308). ACM.

[947] Choi, J. & Sarabandi, K. (2016, June). Non-foster impedance matching of an extremely small folded-dipole antenna at low-VHF band. In *Antennas and Propagation (APSURSI), 2016 IEEE International Symposium on* (pp. 1695-1696). IEEE.

[948] Choi, J., and Sarabandi, K. (2016). Non-foster impedance matching of an extremely small folded dipole antenna at low VHF band. In *IEEE URSI/AP Conference*, Fajardo, Puerto Rico.

[949] Choi, J., and Sarabandi, K. (2016, January). HF/VHF antenna characterization from very-near-field measurements over arbitrary closed surfaces. In *Radio Science Meeting (USNC-URSI NRSM), 2016 United States National Committee of URSI National* (pp. 1-2). IEEE.

[950] Choi, J., and Sarabandi, K. (2016, June). Full-spherical radiation pattern evaluation of low frequency antennas using a novel very-near-field electro-optical system. In *Antennas and Propagation (APSURSI), 2016 IEEE International Symposium on* (pp. 1737-1738). IEEE.

[951] Choi, J., Dagefu, F. T., Sadler, B. M., and Sarabandi, K. (2016). Electrically small folded dipole antenna for HF and low-VHF bands. *IEEE Antennas and Wireless Propagation Letters*, *15*, 718-721.

[952] Choi, J., Sarabandi, K., Dagefu, F. T., and Sadler, B. M. (2016, June). A compact, low-power, low-VHF radio for mobile and wireless communication applications. In *USNC-URSI Radio Science Meeting, 2016* (pp. 71-72). IEEE.

[953] Choudhary, S., Carlone, L., Nieto, C., Rogers, J., Christensen, H. I., and Dellaert, F. (2016, May). Distributed trajectory estimation with privacy and communication constraints: A two-stage distributed Gauss-Seidel approach. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on* (pp. 5261-5268). IEEE.

[954] Choudhary, S., Carlone, L., Nieto, C., Rogers, J., Christensen, H. I., and Dellaert, F. (2016, May). Distributed trajectory estimation with privacy and communication constraints: A two-stage distributed Gauss-Seidel approach. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on* (pp. 5261-5268). IEEE.

[955] Choudhary, S., Carlone, L., Nieto, C., Rogers, J., Liu, Z., Christensen, H. I., and Dellaert, F. (2017). Multi robot object-based SLAM. In *2016 International Symposium on Experimental Robotics*. Springer International Publishing.

[956] Choudhary, S., Carlone, L., Nieto, C., Rogers, J., Liu, Z., Christensen, H. I., and Dellaert, F. (2016) Multi robot object-based SLAM. In *Intl. Symp. on Experimental Robotics*,(Tokyo, JP, Oct 2016. IFRR.

[957] Christensen, D. L., Suresh, S. A., Hahm, K., and Cutkosky, M. R. (2016). Let's All Pull Together: Principles for sharing large loads in microrobot teams. *IEEE Robotics and Automation Letters*, *1*(2), 1089-1096.

[958] Concha, A., Loianno, G., Kumar, V., and Civera, J. (2016, May). Visual-inertial direct SLAM. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on* (pp. 1331-1338). IEEE.

[959] Conroy, J., Nothwang, W., and Gremillion, G. (2016, September). Continuous time rate gyro calibration and motion capture system misalignment estimation using a nonlinear observer. In *Multisensor Fusion and Integration for Intelligent Systems (MFI), 2016 IEEE International Conference on* (pp. 487-492). IEEE.

[960] Corah, M. & Michael, N., (2016) Coordinated active perception and manipulation of a payload with unknown center-of-mass location. In *RSS Workshop on Robot-Environment Interaction for Perception and Manipulation*, Ann Arbor, MI, Jun. 2016.

[961] Cruz, P. J., and Fierro, R. (2015, December). Building coalitions of heterogeneous agents using weighted bipartite graphs. In *Decision and Control (CDC), 2015 IEEE 54th Annual Conference on* (pp. 2822-2828). IEEE.

[962] Cruz, P. J., Sadler, B. M., and Fierro, R. (2016, July). Sensor localization using hybrid RF/optical wireless communications for an aerial data mule. In *American Control Conference (ACC), 2016* (pp. 7085-7091). IEEE.

[963] Cruz, P. J., Sadler, B. M., and Fierro, R. (2016, July). Sensor localization using hybrid RF/optical wireless communications for an aerial data mule. In *American Control Conference (ACC), 2016* (pp. 7085-7091). IEEE.

[964] Cruz, P., Hintz, C., West, J., and Fierro, R., (2016) Optical wireless communication for heterogeneous DARS. 13th Int. Symposium on Distributed Autonomous Robotic Systems (DARS), London, UK, Nov. 6-9, 2016. Digital Proceedings.

[965] Dagefu, F., Verma, G., Choi, J., and Sadler, B. M. (2016, June). Empirical and simulation based analysis of propagation mechanisms at low frequencies. In *Antennas and Propagation (APSURSI), 2016 IEEE International Symposium on* (pp. 1695-1696). IEEE.

[966] Dagefu, F., Verma, G., Choi, J., Sadler, B., and Sarabandi, K. (2016) Empirical and simulation based analysis of propagation mechanisms at low frequencies. In *2016 IEEE International Symposium on Antennas and Propagation/USNC-URSI National Radio Science meeting*, June 2016.

[967] Dai, J., Faraji, H., Gong, C., Hatton, R. L., Goldman, D. I., and Choset, H. (2016) Geometric swimming on a granular surface. *Robotics: Science & Systems (RSS)*.

[968] Desaraju, V. R., and Michael, N. (2016). Experience-driven predictive control. *Robot Learning and Planning (RLP 2016)*, *29*.

[969] Desaraju, V. R., and Michael, N. (2016, May). Fast nonlinear model predictive control via partial enumeration. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on* (pp. 1243-1248). IEEE.

[970] Estrada, M. A., Hockman, B., Bylard, A., Hawkes, E. W., Cutkosky, M. R., and Pavone, M. (2016, May). Free-flyer acquisition of spinning objects with gecko-inspired adhesives. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on* (pp. 4907-4913). IEEE.

[971] Gao, H., and Wei, M. (2016). Domain decomposition in POD-Galerkin projection for flows with moving boundary. In *54th AIAA Aerospace Sciences Meeting* (p. 1102).

[972] Gao, H., and Wei, M. (2016, November). Model order reduction for fluid dynamics with moving solid boundary. In *APS Division of Fluid Dynamics Meeting Abstracts*.

[973] Gil, S., Kumar, S., Mazumder, M., Katabi, D., and Rus, D. (2015). Guaranteeing spoof-resilient multi-robot networks.

[974] Goldberg, J. D., and Fearing, R. S. (2015, September). Force sensing shell using a planar sensor for miniature legged robots. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on* (pp. 1494-1500). IEEE.

[975] Gong, C., Goldman, D. I., and Choset, H. (2016) Simplifying gait design via shape basis optimization. *Robotics: Science & Systems (RSS).*

[976] Gravish, N., Gold, G., Zangwill, A., Goodisman, M. A., and Goldman, D. I. (2015). Glass-like dynamics in confined and congested ant traffic. *Soft Matter*, *11*(33), 6552-6561.

[977] Gremillion, G. M., and Humbert, J. S. (2016). Disturbance rejection with distributed acceleration sensing for small unmanned aircraft systems. *AIAA Journal*, (0), 2233-2246.

[978] Gremillion, G., and Humbert, J. S. (2015, May). Bio-inspired sensing and control for disturbance rejection and stabilization. In *SPIE Defense+ Security* (pp. 94671O-94671O). International Society for Optics and Photonics.

[979] Haldane, D. W., Plecnik, M. M., Yim, J. K., and Fearing, R. S. (2016). Robotic vertical jumping agility via series-elastic power modulation. *Science Robotics*, *1*(1), eaag2048.

[980] Haldane, D. W., Plecnik, M., Yim, J. K., and Fearing, R. S. (2016, October). A power modulating leg mechanism for monopedal hopping. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on* (pp. 4757-4764). IEEE.

[981] Hawkes, E. W., Christensen, D. L., Pope, M. T., and Cutkosky, M. R. (2016). One motor, two degrees of freedom through dynamic response switching. *IEEE Robotics and Automation Letters*, *1*(2), 969-975.

[982] Hrishikeshavan, V., Benedict, M., and Chopra, I. (2014). Identification of flight dynamics of a cylcocopter micro air vehicle in hover. *Journal of Aircraft*, *52*(1), 116-129.

[983] Jafari, M., and Rais-Zadeh, M. (2016, January). A 1550 NM phase change electro-optical shutter. In *Micro Electro Mechanical Systems (MEMS), 2016 IEEE 29th International Conference on* (pp. 655-658). IEEE.

[984] Jam, A., and Sarabandi, K. (2016, June). A horizontally polarized beam-steerable antenna for sub-millimeter-wave polarimetrie imaging and collision avoidance radars. In *Antennas and Propagation (APSURSI), 2016 IEEE International Symposium on* (pp. 789-790). IEEE.

[985] Jam, A., and Sarabandi, K. (2017). A submillimeter-wave near-field measurement setup for on-wafer pattern and gain characterization of antennas and arrays. *IEEE Transactions on Instrumentation and Measurement*, *66*(4), 802-811.

[986] Jayaram, K., and Full, R. J. (2016). Cockroaches traverse crevices, crawl rapidly in confined spaces, and inspire a soft, legged robot. *Proceedings of the National Academy of Sciences*, *113*(8), E950-E957.

[987] Jayaram, K., and Full, R. J. (2016, March). Cockroaches squeezing through crevices. *Integrative and Comparative Biology*, *56*, E102-E102. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[988] Jeon, D., Dong, Q., Kim, Y., Wang, X., Chen, S., Yu, H., and Sylvester, D. (2017). A 23-mW face recognition processor with mostly-read 5T memory in 40-nm CMOS. *IEEE Journal of Solid-State Circuits*.

[989] Kahn, G., Zhang, T., Levine, S., and Abbeel, P. (2016). PLATO: Policy Learning using Adaptive Trajectory Optimization. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Singapore, May 2017.

[990] Karydis, K., Kannappan, P., Tanner, H. G., Jardine, A., and Heinz, J. (2016). Resilience through learning in multi-agent cyber-physical systems. *Frontiers in Robotics and AI*, *3*, 36.

[991] Karydis, K., Stager, A., Tanner, H., and Poulakakis, I. (2016). Experimental validation of a template for navigation of miniature legged robots. In *Int. Symposium on Experimental Robotics*.

[992] Kellen, A., and Benedict, M., (2016) Design, development and performance measurements of a UAV-scale cycloidal rotor. In *Proceedings of the 2016 Annual AIAA Region IV Student Conference*, University of Texas at Arlington, Arlington, TX.

[993] Koppel, A., Sadler, B. M., and Ribeiro, A. (2016, December). Decentralized online optimization with heterogeneous data sources. In *Signal and Information Processing (GlobalSIP), 2016 IEEE Global Conference on* (pp. 515-519). IEEE.

[994] Koppel, A., Sadler, B. M., and Ribeiro, A. (2016, March). Proximity without consensus in online multi-agent optimization. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on* (pp. 3726-3730). IEEE.

[995] Lankford, J., Mayo, D., and Chopra, I. (2016). Computational investigation of insect-based flapping wings for micro air vehicle applications. *International Journal of Micro Air Vehicles*, *8*(2), 64-78.

[996] Lee, J. S., Fearing, R. S., and CHO, K. (2016, August). Compound foot for increased millirobot jumping ability. In *Advances in Cooperative Robotics: Proceedings of the 19th International Conference on Clawar 2016* (p. 71). World Scientific.

[997] Leonardos, S., and Daniilidis, K. (2016). A game-theoretic approach to robust fusion and Kalman filtering under unknown correlations. *arXiv preprint arXiv:1610.01045*.

[998] Levine, S., Finn, C., Darrell, T., and Abbeel, P. (2015). End-to-end training of deep visuomotor policies. *arXiv preprint arXiv:1504.00702*.

[999] Li, C., Kessens, C. C., Young, A., Fearing, R. S., and Full, R. J. (2016, October). Cockroach-inspired winged robot reveals principles of ground-based dynamic self-righting. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on* (pp. 2128-2134). IEEE.

[1000] Li, C., Tian, R., Porter, W., Hammond, Z., Strachan-Olson, D., Kooker, A. W., and Full, R. J. (2016, March). Cockroach-inspired self-righting robots. *Integrative and Comparative Biology*, *56*, E323-E323. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[1001] Libby, T., Edgerly, J. S., and Full, R. J. (2016, March). Clumsy dynamics of rapid backwards running in tube-dwelling webspinners.

Integrative and Comparative Biology, *56*, E127-E127. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[1002] Libby, T., Johnson, A. M., Chang-Siu, E., Full, R. J., and Koditschek, D. E. (2016). Comparative design, scaling, and control of appendages for inertial reorientation. *IEEE Transactions on Robotics*, *32*(6), 1380-1398.

[1003] Linevich, V., Monaenkova, D., and Goldman, D. I. (2016). Robophysical study of excavation in confined environments. *Artificial Life and Robotics*, *21*(4), 460-465.

[1004] Liu, L., and Michael, N. (2016, May). An MDP-based approximation method for goal constrained multi-MAV planning under action uncertainty. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on* (pp. 56-62). IEEE.

[1005] Loianno, G., Mulgaonkar, Y., Brunner, C., Ahuja, D., Ramanandan, A., Chari, M., and Kumar, V. (2016, October). A swarm of flying smartphones. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on* (pp. 1681-1688). IEEE.

[1006] Loianno, G., Watterson, M., and Kumar, V. (2016, May). Visual inertial odometry for quadrotors on SE (3). In *Robotics and Automation (ICRA), 2016 IEEE International Conference on* (pp. 1544-1551). IEEE.

[1007] M'Balé, K., and Josyula, D. (2016). Using automatic case generation to enable advanced behaviors in agents. *Procedia Computer Science*, *88*, 444-449.

[1008] MacFarlane, M., and Humbert, J. (2015). Lagrangian flow structures around a flapping wing. In *AIAA Atmospheric Flight Mechanics Conference* (p. 0012).

[1009] Mathis, A., Donavanik, D., Nothwang, W. (2016) Computationally Efficient Scene Categorization in Complex Dynamic Environments. 2016 IEEE Applied Imagery Pattern Recognition Workshop, Washington D.C., October 10-12, 2016.

[1010] Matthies, L. (2015) Progress toward autonomous aerial mobility. In *Workshop on Vision-based Control and Navigation of Small, Lightweight UAVs*, held in conjunction with *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Hamburg, Germany

[1011] McInroe, B., Astley, H. C., Gong, C., Kawano, S. M., Schiebel, P. E., Rieser, J. M., and Goldman, D. I. (2016). Tail use improves soft substrate performance in models of early vertebrate land locomotors.

[1012] Miao, Z., Thakur, D., Erwin, R. S., Pierre, J., Wang, Y., and Fierro, R. (2016, December). Orthogonal vector field-based control for a multi-robot system circumnavigating a moving target in 3D. In *Decision and Control (CDC), 2016 IEEE 55th Conference on* (pp. 6004-6009). IEEE.

[1013] Miao, Z., Wang, Y., and Fierro, R. (2015, December). Collision-free consensus via order preservation in multi-agent networks. In *Decision and Control (CDC), 2015 IEEE 54th Annual Conference on* (pp. 2999-3005). IEEE.

[1014] Miao, Z., Wang, Y., and Fierro, R. (2016). Collision-free consensus in multi-agent networks: A monotone systems perspective. *Automatica*, *64*, 217-225.

[1015] Moosavifar, M., Ansari, A., and Rais-Zadeh, M. (2016, October). An AlN-on-Si resonant IR sensor array with a large temperature coefficient of frequency. In *Sensors, 2016 IEEE* (pp. 1-3). IEEE.

[1016] Murali, V., Nieto, C., Choudhary, S., and Christensen, H. I. (2016, October). Active planning based extrinsic calibration of exteroceptive sensors in unknown environments. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on* (pp. 2498-2505). IEEE.

[1017] Ng, P. Q., Springthorpe, D., and Full, R. J. (2016, March). Mobile manipulation: Ghost crab climbing using pincer-like claws. *Integrative and Comparative Biology*, *56*, E161-E161. Journals Dept, 2001 Evans Rd, Cary, NC: Oxford Univ Press Inc.

[1018] Özaslan, T., Mohta, K., Keller, J., Mulgaonkar, Y., Taylor, C. J., Kumar, V., and Hood, T. (2016, October). Towards fully autonomous visual inspection of dark featureless dam penstocks using MAVs. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on* (pp. 4998-5005). IEEE.

[1019] Pope, M. T., and Cutkosky, M. R. (2016, July). Thrust-assisted perching and climbing for a bioinspired UAV. In *Conference on Biomimetic and Biohybrid Systems* (pp. 288-296). Springer International Publishing.

[1020] Qian, F., Zhang, T., Korff, W., Umbanhowar, P. B., Full, R. J., and Goldman, D. I. (2015). Principles of appendage design in robots and animals determining terradynamic performance on flowable ground. *Bioinspiration & biomimetics*, *10*(5), 056014.

[1021] Runco, C. C., Benedict, M., and Coleman, D. (2016). Design and development of a meso-scale cyclocopter. In *AIAA Infotech@ Aerospace* (p. 1255).

[1022] Runco, C., Coleman, D., and Benedict, M., (2016) Development of the world's smallest cyclocopter. In *Proceedings of the 72nd Annual National Forum of the American Helicopter Society*, West Palm Beach, FL.

[1023] Sadeghi, M. M., Kim, H. S., Peterson, R. L. B., and Najafi, K. (2016). Electrostatic micro-hydraulic systems. *J. Microelectromechanical Systems*, *25*(3), 557-569.

[1024] Savoie, W., Pazouki, A., Negrut, D., and Goldman, D. (2015). Smarticles: design and construction of smart particles to aid discovery of principles of smart, active granular matter. *SWARM 2015* conference paper.

[1025] Shariati, A., Mohta, K., and Taylor, C. J. (2016, October). Recovering relative orientation and scale from visual odometry and ranging radio measurements. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on* (pp. 3627-3633). IEEE.

[1026] Shrestha, E., Hrishikeshavan, V., Benedict, M., Yeo, D., and Chopra, I. (2016). Development of control strategies for a twin-cyclocopter in forward flight. *Journal of the American Helicopter Society*, *61*(4), 1-9.

[1027] Shrestha, E., Hrishikeshavan, V., Yeo, D., and Chopra, I., (2015). Flight dynamics modeling and system identification of a cyclocopter in forward flight. In *71st Annual Forum of the American Helicopter Society*, Virginia Beach, VA.

[1028] Shrestha, E., Martz, V., Chopra, I., and Benedict, M., (2015). Design and hover testing of a 50-gram meso-cyclocopter. In *AHS International Specialists' Meeting on Unmanned Rotorcraft and Network Centric Operations*, Chandler, AZ

[1029] Shrestha, E., Yeo, D., Hrishikeshavan, V., Benedict, M., and Chopra, I., (2016). Gust disturbance rejection study of a cyclocopter micro air vehicle. In *Proceedings of the 72nd Annual Forum of the American Helicopter Society*, West Palm Beach, FL.

[1030] Stager, A., and Tanner, H. G. (2016, December). Stochastic behavior of robots that navigate by interacting with their environment. In *Decision and Control (CDC), 2016 IEEE 55th Conference on* (pp. 6871-6876). IEEE.

[1031] Stelmakh, V., Chan, W. R., Ghebrebrhan, M., Senkevich, J., Joannopoulos, J. D., Soljačić, M., and Celanović, I. (2016). Sputtered tantalum photonic crystal coatings for high-temperature energy conversion applications. *IEEE Transactions on Nanotechnology*, *15*(2), 303-309.

[1032] Stelmakh, V., Chan, W. R., Ghebrebrhan, M., Soljacic, M., Joannopoulos, J. D., and Celanovic, I. (2015). Photonic crystal emitters for thermophotovoltaic energy conversion. In *Journal of Physics: Conference Series* (Vol. 660, No. 1, p. 012080). IOP Publishing.

[1033] Stelmakh, V., Chan, W. R., Ghebrebrhan, M., Soljacic, M., Joannopoulos, J. D., and Celanovic, I. (2016, November). Fabrication of an omnidirectional 2d photonic crystal emitter for thermophotovoltaics. *Journal of Physics: Conference Series*, *773*(1), 012037. IOP Publishing.

[1034] Stelmakh, V., Chan, W. R., Joannopoulos, J. D., Soljacic, M., Celanovic, I., and Sablon, K. (2016, January). Improved thermal emitters for thermophotovoltaic energy conversion. In *ASME 2016 5th International*

Conference on Micro/Nanoscale Heat and Mass Transfer (pp. V001T05A015-V001T05A015). American Society of Mechanical Engineers.

[1035] Sun, J., and Tanner, H. G. (2015). Constrained decision-making for low-count radiation detection by mobile sensors. *Autonomous Robots*, *39*(4), 519-536.

[1036] Tabandeh, M., Wei, M., and Collins, J. P. (2016). On the symmetrization in POD-Galerkin model for linearized compressible flows. In *54th AIAA Aerospace Sciences Meeting* (p. 1106).

[1037] Tabib, W., Corah, M., Michael, N., and Whittaker, R. (2016, October). Computationally efficient information-theoretic exploration of pits and caves. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on* (pp. 3722-3727). IEEE.

[1038] Tang, Y., and Najafi, K. (2015, December). A two-gap capacitive structure for high aspect-ratio capacitive sensor arrays. In *Electron Devices Meeting (IEDM), 2015 IEEE International* (pp. 18-2). IEEE.

[1039] Tang, Y., and Najafi, K. (2016, February). High aspect-ratio low-noise multi-axis accelerometers made from thick silicon. In *Inertial Sensors and Systems, 2016 IEEE International Symposium on* (pp. 121-124). IEEE.

[1040] Thomas, J., Loianno, G., Daniilidis, K., and Kumar, V. (2016). Visual servoing of quadrotors for perching by hanging from cylindrical objects. *IEEE Robotics and Automation Letters*, *1*(1), 57-64.

[1041] Thomas, J., Loianno, G., Daniilidis, K., and Kumar, V. (2016, May). The role of vision in perching and grasping for MAVs. In *SPIE Defense+ Security* (pp. 98361S-98361S). International Society for Optics and Photonics.

[1042] Thomas, J., Loianno, G., Pope, M., Hawkes, E. W., Estrada, M. A., Jiang, H., and Kumar, V. (2015, August). Planning and control of aggressive maneuvers for perching on inclined and vertical surfaces. In *ASME 2015 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (pp. V05CT08A012-V05CT08A012). American Society of Mechanical Engineers.

[1043] Thomas, J., Pope, M., Loianno, G., Hawkes, E. W., Estrada, M. A., Jiang, H., and Kumar, V. (2016). Aggressive flight with quadrotors for perching on inclined surfaces. *Journal of Mechanisms and Robotics*, *8*(5), 051007.

[1044] Tokekar, P., and Kumar, V. (2015, September). Visibility-based persistent monitoring with robot teams. In *Intelligent Robots and Systems*

(IROS), 2015 IEEE/RSJ International Conference on (pp. 3387-3394). IEEE.

[1045] Travers, M., Whitman, J., Schiebel, P., Goldman, D., and Choset, H. (2016). Shape-based compliance in locomotion. In *Robotics: Science and Systems (RSS)* (No. 5, p. 10).

[1046] Tron, R., Thomas, J., Loianno, G., Daniilidis, K., and Kumar, V. (2016). A distributed optimization framework for localization and formation control: applications to vision-based measurements. *IEEE Control Systems*, *36*(4), 22-44.

[1047] Tron, R., Thomas, J., Loianno, G., Daniilidis, K., and Kumar, V. (2016, December). Bearing-only formation control with auxiliary distance measurements, leaders, and collision avoidance. In *Decision and Control (CDC), 2016 IEEE 55th Conference on* (pp. 1806-1813). IEEE.

[1048] Verma, G., Dagefu, F., Sadler, B. M., and Sarabandi, K. (2016, June). Direction of arrival estimation with the received signal strength gradient at the lower VHF band. In *Antennas and Propagation (APSURSI), 2016 IEEE International Symposium on* (pp. 1695-1696). IEEE.

[1049] Wang, S., Jiang, H., and Cutkosky, M. R. (2016, October). A palm for a rock climbing robot based on dense arrays of micro-spines. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on* (pp. 52-59). IEEE.

[1050] Wang, X., Chan, W. R., Stelmakh, V., Soljacic, M., Joannopoulos, J. D., Celanovic, I., and Fisher, P. H. (2015). Prototype of radioisotope thermophotovoltaic system using photonic crystal spectral control. *Journal of Physics: Conference Series*, *660*(1), 012034. IOP Publishing.

[1051] Wang, X., Chan, W., Stelmakh, V., Celanovic, I., and Fisher, P. (2016). Toward high performance radioisotope thermophotovoltaic systems using spectral control. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, *838*, 28-32.

[1052] Wei, M., and Xu, M. (2015, November). Adjoint-based optimization for the understanding of the aerodynamics of a flapping plate. In *APS Meeting Abstracts*.

[1053] Wu, X. A., Huh, T. M., Mukherjee, R., and Cutkosky, M. (2016). Integrated Ground Reaction Force Sensing and Terrain Classification for Small Legged Robots. *IEEE Robotics and Automation Letters*, *1*(2), 1125-1132.

[1054] Xiang, J., Li, Z., Blaauw, D., Kim, H. S., and Chakrabarti, C. (2016, September). Low complexity optical flow using neighbor-guided semi-

global matching. In *Image Processing (ICIP), 2016 IEEE International Conference on* (pp. 4483-4487). IEEE.

[1055] Xu, M., and Wei, M. (2016). Using adjoint-based optimization to study kinematics and deformation of flapping wings. *Journal of Fluid Mechanics*, *799*, 56-99.

[1056] Xu, M., Wei, M., Yang, T., and Lee, Y. S. (2016). An embedded boundary approach for the simulation of a flexible flapping wing at different density ratio. *European Journal of Mechanics-B/Fluids*, *55*, 146-156.

[1057] Yao, J. W., Desaraju, V. R., and Michael, N. (2016, October). Experience-based models of surface proximal aerial robot flight performance in wind. In *International Symposium on Experimental Robotics* (pp. 563-573). Springer, Cham.

[1058] Yektakhah, B., and Sarabandi, K. (2016) A low profile wideband circularly polarized omnidirectional antenna using an array of sectorial loop antennas. In *IEEE Antennas and Propagation Society International Symposium (APS/URSI)*.

[1059] Yektakhah, B., and Sarabandi, K. (2016). All-directions through-the-wall radar imaging using a small number of moving transceivers. *IEEE Transactions on Geoscience and Remote Sensing*, *54*(11), 6415-6428.

[1060] Zhang, T., Kahn, G., Levine, S., and Abbeel, P. (2016, May). Learning deep control policies for autonomous aerial vehicles with mpc-guided policy search. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on* (pp. 528-535). IEEE.

[1061] Zhang, W., and Wei, M. (2016, November). A low-order Galerkin model based on DMD and adjoint-DMD modes. In *APS Division of Fluid Dynamics Meeting Abstracts*.

FY17

[1062] Badrya, C., Sridharan, A., Baeder, J. D., and Kroninger, C. M. (2017). Multi-Fidelity Coupled Trim Analysis of a Flapping-Wing Micro Air Vehicle Flight. *Journal of Aircraft*, *54*(5), 1614-1630.

[1063] Choi, J., Sarabandi, K., Dagefu, F. T., and Sadler, B. M. (2017, July). A Non-Foster matched dipole for a low-vhf mobile transmitter system. In *Antennas and Propagation & USNC/URSI National Radio Science Meeting, 2017 IEEE International Symposium on* (pp. 2357-2358). IEEE.

[1064] Choudhary, S., Carlone, L., Nieto, C., Rogers, J., Christensen, H. I., and Dellaert, F. (2017). Distributed mapping with privacy and communication constraints: Lightweight algorithms and object-based models. *The International Journal of Robotics Research*, *36*(12), 1286-1311.

[1065] Cigla, C., Brockers, R., and Matthies, L. (2017, July). Image-based visual perception and representation for collision avoidance. In *IEEE International Conference on Computer Vision and Pattern Recognition, Embedded Vision Workshop*.

[1066] Cigla, C., Brockers, R., and Matthies, L. (2017, March). Gaussian mixture models for temporal depth fusion. In *Applications of Computer Vision (WACV), 2017 IEEE Winter Conference on* (pp. 889-897). IEEE.

[1067] Corah, M., and Michael, N. (2017). Efficient online multi-robot exploration via distributed sequential greedy assignment. In *Proceedings of robotics: science and systems*.

[1068] Corah, M., and Michael, N. (2017, May). Active estimation of mass properties for safe cooperative lifting. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on* (pp. 4582-4587). IEEE.

[1069] Cruz, P. J., and Fierro, R. (2017). Cable-suspended load lifting by a quadrotor UAV: hybrid model, trajectory generation, and control. *Autonomous Robots*, *41*(8), 1629- 1643.

[1070] Dagefu, F. T., Verma, G., Sadler, B. M., Kozick, R., and Sarabandi, K. (2017, July). Full-wave analysis of time of arrival based localization with polarization diversity. In *USNC-URSI Radio Science Meeting (Joint with AP-S Symposium), 2017* (pp. 5-6). IEEE.

[1071] Deetjen, M. E., Biewener, A. A., and Lentink, D. (2017). High-speed surface reconstruction of a flying bird using structured light. *Journal of Experimental Biology*, *220*(11), 1956-1961.

[1072] Deetjen, M., and Lentink, D. (2017). High-Speed Surface Reconstruction of Flying Birds Using Structured Light. *Bulletin of the American Physical Society*, *62*.

[1073] Desaraju, V. R., and Michael, N. (2017, May). Leveraging experience for computationally efficient adaptive nonlinear model predictive control. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on* (pp. 5314-5320). IEEE.

[1074] Desaraju, V. R., Spitzer, A. E., and Michael, N. (2017, July). Experience-driven Predictive Control with Robust Constraint Satisfaction under Time-Varying State Uncertainty. In *Robotics: Science and Systems Conference (RSS), July*.

[1075] Faust, A., Palunko, I., Cruz, P., Fierro, R., and Tapia, L. (2017). Automated aerial suspended cargo delivery through reinforcement learning. *Artificial Intelligence*, *247*, 381-398.

[1076] Gil, S., Kumar, S., Mazumder, M., Katabi, D., and Rus, D. (2017). Guaranteeing spoof- resilient multi-robot networks. *Autonomous Robots*, *41*(6), 1383-1400.

[1077] Gurel, N. Z. (2016). *Frequency domain characterization of optic flow and vision- based ocellar sensing for rotational motion* (Doctoral dissertation).

[1078] Himmelberg, B., and Benedict, M. (2017). Performance Measurements of Meso-Scale Cyclodial Rotors in Hover. In *55th AIAA Aerospace Sciences Meeting* (p. 2009).

[1079] Hintz, C., Ahmad, S., Kloeppel, J., and Fierro, R. (2017, August). Robust hybrid control for swinging-up and balancing an inverted pendulum attached to a UAV. In *Control Technology and Applications (CCTA), 2017 IEEE Conference on* (pp. 1550- 1555). IEEE.

[1080] Hunt, N. H., Frendberg-Mates, E., Jinn, J., Robin, A., Jacobs, L. F., and Full, R. J. (2017, March). Squirrels Running on Compliant Branches: When to Leap?. In *Integrative and Comparative Biology* (Vol. 57, pp. E299-E299). Journals Dept, 2001 Evans Rd, Cary, NC 27513 USA: Oxford Univ Press Inc.

[1081] Jam, A., and Sarabandi, K. (2017). A Submillimeter-Wave Near-Field Measurement Setup for On-Wafer Pattern and Gain Characterization of Antennas and Arrays. *IEEE Transactions on Instrumentation and Measurement*, *66*(4), 802-811.

[1082] Jam, A., East, J., and Sarabandi, K. (2017, June). A micromachined packaging with incorporated RF-choke for integration of active chips at submillimeter-wave frequencies. In *Microwave Symposium (IMS), 2017 IEEE MTT-S International* (pp. 507-509). IEEE.

[1083] Kroninger C, Axial Propulsion with Flapping & Rotating Wings - Minimum Power from Hover through Cruise. *Bioinspiration and Biomimetics.13(3),*036012.

[1084] Lankford, J. L. and Chopra, I.:"Coupled CFD/CSD Study of a Flexible MAV-Scale Flapping Wing in the Hover Flight Condition"Presented at the AHS 73rd Annual Forum in Fort Worth, Texas, May 2017.

[1085] Lee, A. X., Levine, S., and Abbeel, P. (2017). Learning visual servoing with deep features and fitted q-iteration. *arXiv preprint arXiv:1703.11000*.

[1086] Lee, J. S., Fearing, R. S., Cho, K. J., and Full, R. J. (2017, March). Crickets Jumping from Diverse Substrates Inspire Leg Design in a Millirobot. In *Integrative and Comparative Biology* (Vol. 57, Pp. E324-E324). Journals Dept, 2001 Evans Rd, Cary, NC 27513 USA: Oxford Univ Press Inc.

[1087] Lee, K., Park, S., Park, S. Y., Cho, J., and Yoon, E. (2017, June). A 272.49 pJ/pixel CMOS image sensor with embedded object detection and bio-inspired 2D optic flow generation for nano-air-vehicle navigation. In *VLSI Circuits, 2017 Symposium on* (pp. C294-C295). IEEE.

[1088] Leonardos, S., and Daniilidis, K. (2017, May). A game-theoretic approach to Robust Fusion and Kalman filtering under unknown correlations. In *American Control Conference (ACC), 2017* (pp. 2568-2573). IEEE.

[1089] Leonardos, S., Zhou, X., and Daniilidis, K. (2017, May). Distributed consistent data association via permutation synchronization. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on* (pp. 2645-2652). IEEE.

[1090] Li, C., Kessens, C. C., Fearing, R. S., and Full, R. J. (2017). Mechanical principles of dynamic terrestrial self-righting using wings. *Advanced Robotics*, *31*(17), 881-900.

[1091] Libby, T., and Full, R. J. (2017, March). Variable Limb Function Results in Similar Turning Behavior in Lizards. In *Integrative And Comparative Biology* (Vol. 57, Pp. E328-E328). Journals Dept, 2001 Evans Rd, Cary, NC 27513 USA: Oxford Univ Press Inc.

[1092] Loianno, G., Brunner, C., McGrath, G., and Kumar, V. (2017). Estimation, control, and planning for aggressive flight with a small quadrotor with a single camera and IMU. *IEEE Robotics and Automation Letters*, *2*(2), 404-411.

[1093] Mathis, A., Donavanik, D., and Nothwang, W. (2016, October). Computationally efficient scene categorization in complex dynamic environments. In *Applied Imagery Pattern Recognition Workshop (AIPR), 2016 IEEE* (pp. 1-6). IEEE.

[1094] Miao, Z., Wang, Y., and Fierro, R. (2017). Cooperative circumnavigation of a moving target with multiple nonholonomic robots using backstepping design. Systems & Control Letters, 103, 58-65.

[1095] Nogar, S., and Kroninger, C. (2018). Development of a Hybrid Micro Air Vehicle Capable of Controlled Transition. IEEE Robotics and Automation Letters.

[1096] Özaslan, T., Loianno, G., Keller, J., Taylor, C. J., Kumar, V., Wozencraft, J. M., and Hood, T. (2017). Autonomous navigation and mapping for inspection of penstocks and tunnels with MAVs. IEEE Robotics and Automation Letters, 2(3), 1740-1747.

[1097] Rezaian, E., and Wei, M. (2017). Obtaining a Stable Galerkin ROM in Presence of Shock-Vortex Interaction. In 55th AIAA Aerospace Sciences Meeting (p. 1008).

[1098] Richter, C., and Roy, N. (2017). Safe visual navigation via deep learning and novelty detection. In Proc. of the Robotics: Science and Systems Conference.

[1099] Roderick, W. R., Jiang, H., Wang, S., Lentink, D., and Cutkosky, M. R. (2017, July). Bioinspired Grippers for Natural Curved Surface Perching. In Conference on Biomimetic and Biohybrid Systems (pp. 604-610). Springer, Cham.

[1100] Runco, C. C., Benedict, M., and Coleman, D. (2016). Design and Development of a Meso-Scale Cyclocopter. In AIAA Infotech@ Aerospace (p. 1255).

[1101] Shamwell, E. J., Nothwang, W. D., and Perlis, D. Deepefference: Learning to predict the sensory consequences of action through deep correspondence. In Development and Learning and Epigenetic Robotics (ICDL), 2017 IEEE International Conference on. IEEE, Accepted.

[1102] Shrestha, E., Hrishikeshavan, V., Benedict, M., Yeo, D., and Chopra, I. (2016). Development of Control Strategies for a Twin-Cyclocopter in Forward Flight. Journal of the American Helicopter Society, 61(4), 1-9.

[1103] Shrestha, E., Yeo, D., Benedict, M., and Chopra, I. (2017). Development of a meso- scale cycloidal-rotor aircraft for micro air vehicle application. International Journal of Micro Air Vehicles, 9(3), 218-231.

[1104] Smith, Z. F., Jones, A. R., and Hrynuk, J. T. (2018). Micro Air Vehicle Scale Gust- Wing Interaction in a Wind Tunnel. In 2018 AIAA Aerospace Sciences Meeting (p. 0573).

[1105] Thomas, J., Welde, J., Loianno, G., Daniilidis, K., and Kumar, V. (2017). Autonomous flight for detection, localization, and tracking of moving targets with a small quadrotor. IEEE Robotics and Automation Letters, 2(3), 1762-1769.

[1106] Tran, J., Gao, H., Sirohi, J., and Wei, M. (2017). Reduced-order methodology for prediction of loads generated by a flexible flapping wing. *International Journal of Micro Air Vehicles*, 1756829317708318.

[1107] Tron, R., Thomas, J., Loianno, G., Daniilidis, K., and Kumar, V. (2016, December). Bearing-only formation control with auxiliary distance measurements, leaders, and collision avoidance. In *Decision and Control (CDC), 2016 IEEE 55th Conference on* (pp. 1806-1813). IEEE.

[1108] Verma, G., Dagefu, F., Sadler, B. M., and Sarabandi, K. (2017, July). Localization via the received signal strength gradient at lower VHF. In *USNC-URSI Radio Science Meeting (Joint with AP-S Symposium), 2017* (pp. 7-8). IEEE.

[1109] Walther, C., Coleman, D., Benedict, M., and Lakshminarayan, V. K., "Experimental and Computational Studies to Understand Unsteady Aerodynamics of Cycloidal Rotors in Hover at Ultra-low Reynolds Numbers," Proceedings of the 73rd Annual National Forum of the American Helicopter Society, Fort Worth, TX, May 9–11, 2017.

List of Symbols, Abbreviations, and Acronyms

2-D	2-dimensional
2.5-D	2.5-dimensional
3-D	3-dimensional
APP	Annual Program Plan
ARL	US Army Research Laboratory
AV	AeroVironment
CFD	computational fluid dynamics
CLARK	Common Light Autonomous Robotic Kit
CNC	Communication, Networking, and Coordination
COTS	commercial off the shelf
CPD	Capability Production Document
CTA	Collaborative Technology Alliance
DARPA	Defense Advanced Research Projects Agency
DEA	dielectric elastomer actuator
DGS	Distributed Gauss-Seidel
DNC	differential neural computer
DOD	Department of Defense
DSP	digital signal processer
EO	electro-optic
EPC	Experience-driven Predictive Control
FLA	Fast Lightweight Autonomy
FOV	field of view
FY	fiscal year
GaAlN	gallium aluminum nitride
GaN	gallium nitride

GPS	global positioning system
IMU	inertial measurement unit
IR	infrared
ISR	intelligence, surveillance, and reconnaissance
JPL	Jet Propulsion Laboratory
JX	Joint Experiments
LSTM	long short term memory
LTI	linear time invariant
LWPR	locally weighted projection regression
MAST	Micro Autonomous Systems and Technology
MAV	micro air vehicle
MBDyn	high-fidelity coupled multibody
MCE	Mobility, Control, and Energetics
MEMS	microelectromechanical systems
MIT	Massachusetts Institute of Technology
NEDT	noise equivalent delta temperature
NLOS	non-line-of-sight
NMPC	nonlinear model predictive control
NPE	nonlinear partial enumeration
PI	principal investigator
PV	photovoltaic
PZT	lead zirconate titanate
QoS	quality of service
RANS	Reynolds averaged Navier Stokes
RAS	Robotic and Autonomous Systems Strategy
RC	radio-controlled
RF	radio frequency

RFT	Resistive Force Theory
SAA	sense-and-avoid
SBS	Soldier-borne sensor
SfM	structure from motion
SLAM	simultaneous localization and mapping
SLIP	spring-loaded inverted pendulum
SPP	Sensing, Perception, and Processing
SRM	short-range micro
STMD	small target motion detection
SUAS	small unmanned aircraft system
SUGV	small unmanned ground vehicle
SWaP	size, weight, and power
TPV	thermophotovoltaics
UAV	unmanned aerial vehicle
UC	University of California
VHF	very high frequency

1 (PDF) DEFENSE TECHNICAL INFORMATION CTR DTIC OCA

2 (PDF) DIR ARL
IMAL HRA
RECORDS MGMT
RDRL DCL
TECH LIB

1 (PDF) GOVT PRINTG OFC
A MALHOTRA

6 (PDF) ARL
RDRL SER L
A MATHIS
B PIEKARSKI
W NOTHWANG
RDRL CII A
D BARAN
RDRL VTA
C KRONINGER
RDRL VT
B SADLER