



AFRL S&T Technical Forecasting Summit

David Hahn
UNIVERSITY OF FLORIDA GAINESVILLE
207 GRINTER HALL
GAINESVILLE, FL 32611-5500

12/07/2018
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/RTB2
Arlington, Virginia 22203
Air Force Material Command

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> OMB No. 0704-0188	
The 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 of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to the Department of Defense, Executive Service Directorate (0704-0188). 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 ORGANIZATION.					
1. REPORT DATE (DD-MM-YYYY) 6/30/18		2. REPORT TYPE Final		3. DATES COVERED (From - To) 3/1/2018-9/28/2018	
4. TITLE AND SUBTITLE (SECAF 2030 S&T) Technical Forecasting Summit				5a. CONTRACT NUMBER FA-9550-18-1-0178	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Dr. David W. Hahn Dr. Warren Dixon				5d. PROJECT NUMBER P0075463, 64, 65, 66	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Florida Division of Sponsored Research 1523 Union Road Rm 207 Gainesville, FL 32611-1941				8. PERFORMING ORGANIZATION REPORT NUMBER 5E687	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research Atlanta Atlanta Regional Office 100 Alabama St., SW, Suite 4R15 Atlanta, GA 30303-3104				10. SPONSOR/MONITOR'S ACRONYM(S) ONRRO	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) N66020	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for Public Release					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The University of Florida conducted a Technical Forecasting Summit, bringing together thought leaders from academia and the private sector to help guide the Air Force 2030 Science and Technology Strategy. Led by UF Mechanical and Aerospace Engineering, the overall team of participants combined a broad depth of multidisciplinary expertise with considerable experience and knowledge relevant to the Air Force domains in air, space, and cyber. The summit focused on a diverse set of technical themes, each championed by a team of domain experts and facilitators. Specific topics addressed during the summit were Advanced Modeling, Advanced and Additive Manufacturing, Excited State Materials Performance, Energetics and Fuels, Hypersonics and High-speed Effects, Cyber Systems, Machine Learning, Autonomy and AI, Sensing and Communication, Space Systems, and Unmanned Aerial Vehicles (UAVs). Through outreach and networking, the two-day summit engaged nearly 100 participants from across the nation, and culminated in the following technical roadmap and technical forecasting report.					
15. SUBJECT TERMS Advanced Modeling, Advanced and Additive Manufacturing, Excited State Materials, Energetics and Fuels, Hypersonics and High Speed effects, Cyber Systems, Unmanned Aerial Vehicles, Space Systems, Sensing and Communication, Machine Learning, Autonomy and AI					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 76	19a. NAME OF RESPONSIBLE PERSON Dr. David W. Hahn
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code) 352-392-0807

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATES COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33615-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report, e.g. AFOSR-82-1234.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. 61101A.

5d. PROJECT NUMBER. Enter all project numbers as they appear in the report, e.g. 1F665702D1257; ILIR.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

UF Forecasting Summit

Draft Report as Submitted under AFOSR Contract # FA9550-18-1-0178

Principle Investigators

David W. Hahn

Warren E. Dixon

ABSTRACT

The University of Florida conducted a Technical Forecasting Summit, bringing together thought leaders from academia and the private sector to help guide the Air Force 2030 Science and Technology strategy. Led by UF Mechanical and Aerospace Engineering, the overall team of participants combined a broad depth of multidisciplinary expertise with considerable experience and knowledge relevant to the Air Force domains in air, space, and cyber. The summit focused on a diverse set of technical themes, each championed by a team of domain experts and facilitators. Specific topics addressed during the summit were Advanced Modeling, Advanced and Additive Manufacturing, Excited State Materials Performance, Energetics and Fuels, Hypersonics and High-speed Effects, Cyber Systems, Machine Learning, Autonomy and AI, Sensing and Communication, Space Systems, and Unmanned Aerial Vehicles (UAVs). Through outreach and networking, the two-day summit engaged nearly 100 participants from across the nation, and culminated in the following technical roadmap and technical forecasting report.

Table of Contents

Advanced Modeling	5
Advanced and Additive Manufacturing	8
Excited State Materials Performance	10
Energetics and Fuels	14
Hypersonics and High-speed Effects	22
Cyber Systems	26
Machine Learning, Autonomy and AI	32
Sensing and Communication	42
Space Systems	48
Unmanned Aerial Vehicles (UAVs)	56
App. I – Summit Details	66
App. II – Summit Topics	67
App. III – List of Participants	68
App. IV – List of References and Bibliography	70

Overview

Secretary of Air Force Heather Wilson has initiated a broad *S&T Strategy 2030* study to help guide the Air Force to meet the challenge of exploiting rapidly developing technologies in the context of air, space and/or cyber domains, including the consideration of peer and near-peer adversaries. Motivated in part by previous S&T forecasting studies, including Dr. von Karman's *Toward New Horizons* study, the goal is to provide a broad blue sky vision of issues and novel ideas relevant to national defense domains, including as relevant a summary of the current state of the art, the identification and framing of challenges and opportunities, and ultimately, to conduct long-range technical forecasting for basic and applied research possibilities. By inviting academic and private industry communities to identify technologies and participate, dual public and private sector insights and perspectives are realized.

In a system of systems context reviewing emerging technology and drawing top-level assessments based on open know of the weapons and communication environment of the state of the Air Force today, we identified a suite of emerging technologies that support the 21st Century challenges. The following narrative covers these ten topics addressed in the S&T summit, namely:

- (1) Advanced Modeling
- (2) Advanced and Additive Manufacturing
- (3) Excited State Materials Performance
- (4) Energetics and Fuels
- (5) Hypersonics and High-speed Effects
- (6) Cyber Systems
- (7) Machine Learning, Autonomy and AI
- (8) Sensing and Communication
- (9) Space Systems
- (10) Unmanned Aerial Vehicles (UAVs)

1. Advanced Modeling

US Air Force technological goals can be achieved only through long-term and steady effort towards the development of advanced models. While these technological goals are diverse, ranging from continuous development of existing capabilities that will ensure that the USAF maintains air dominance to extending and reinforcing the AF capabilities through exploration and development of new technologies, *advanced modeling* is essential to obtaining the greatest effectiveness in a budget-constrained environment. For example, fuel cost savings by increases in turbine engine efficiency, advances in lightweight materials and multifunctional structures, and advanced aerodynamic concepts and technologies can only be achieved if robust and predictive models are available and transferable to AF scientists through implementation into government codes.

At present, decision-making is largely based on extensive experimentation and “trial and error design”, which by nature requires a lot of resources, both human and financial. While this approach has allowed the development of advanced AF capabilities, it cannot keep up with the pace of development of new processes/concepts, and new materials, either because experimental data cannot be acquired for the environments in which these systems need to operate, or because direct experimentation is so expensive that the cost strongly limits the number of tests. Moreover, AF systems are complex multi-material, multi-physics systems involving a large number of parameters, so the “trial and error” method extends the development time, and ultimately, the overall cost of the technology.

While huge cost savings can be achieved by reducing the number of experiments required to develop/verify/validate AF systems, at present there is a lack of reliable and performant software/models that enable virtual testing at the engineering scale. For the development of virtual prototypes it is critical to use models that are truly predictive, meaning that the models should be able to predict the thermo-mechanical behavior of the AF systems and their surrounding environment for scenarios that cannot be mimicked on the ground. It is to be noted that in order to be truly predictive, all the models should be valid for general loadings i.e. in the six-dimensional space (i.e. including loadings corresponding to three normal stresses and three shear stresses, or the three translations and three rotations for kinematics/dynamics modeling). However, it is well known that for lower length scales there is a lack of stress-strain data even under simple loading scenarios.

The overall goal of advanced modeling is to establish the set of constitutive equations describing the thermomechanical behavior of materials and systems and numerical capabilities to enable the use of materials or technologies for AF applications. To be transformative and reach the AF technological goals, the advanced modeling field has to overcome many challenges. The novel materials (e.g. new lightweight materials or extreme environment materials) that are being developed push the limit of applicability/validity of existing constitutive models. The transition between conventional material systems to lightweight materials can be achieved only if the set of

constitutive equations describing the thermo-mechanical behavior of the novel materials are truly predictive and can be used in a computational framework allowing virtual prototyping. While new manufacturing processes, such as additive manufacturing or near shape-forming, are very promising for the development of future AF systems and applications, the reliability and consistency of structures made using these new manufacturing processes can only be achieved through modeling. However, to fully take advantage of these new processes, and fully exploit their potential, it is vital that new constitutive models are developed.

With the emergence of complex systems for extreme environments (high temperature, high-velocity, high strain-rate), it is imperative that the modeling framework should account for multi-physics couplings. While for conventional systems and designs it is possible to neglect some of the couplings, for the development of advanced AF systems, the modeling framework and the set of constitutive equations describing each physical system should be rethought. As an example, the challenges of atmospheric hypersonic flight can be overcome only through a new theoretical and computational modeling framework. Due to the harsh environments, especially the associated rise in temperature, the structure will undergo irreversible (permanent) deformation leading to critical changes in geometry. Accounting for these plastic deformations will affect both the fluid dynamics analysis and the thermal analysis, and vice versa. Therefore, for atmospheric hypersonic systems design there is a strong need to develop a new multi-physics modeling framework which couples Plasticity-Computational Fluid Dynamics (CFD)-Thermal analysis.

In order to have the desired impact on future AF systems and applications, emphasis should be placed on capturing the complex physics while ensuring that the models developed are easy to transition. For this purpose, more emphasis should be put on development of rational strategies for parameters identification, and applicability to large scale problems.

It is worth noting that although advanced modeling is a separate technical emphasis area in the S&T strategy developed in this report, it is in fact essential to overcoming all the technological challenges to future AF dominance. This is due to the fact that all future AF systems will involve multi-materials of increased structural complexity and interactions. On one hand, modeling will advance the fundamental understanding of the thermomechanical behavior of materials and will be used to reverse-engineer and design materials and systems with tailored properties/architecture; on the other hand, truly predictive models of the thermo-mechanical behavior will extend the design space and lead to innovative and transformative design solutions. Ultimately, one of the most outstanding benefits of a long-term and steady effort toward the development of advanced models is the cost savings resulting from more efficient technologies and the use of virtual prototyping in the design loop, thus reducing the number of full-scale experiments needed prior to integration of new materials and concepts in AF systems. Clearly, advanced modeling is the key enabler for reaching the technological goals in all the S&T areas identified in this report. In the following, more specific fundamental research challenges/needs are presented.

Constitutive Modeling

USAF has a long history of enabling basic research aimed at development of mathematical methods and algorithms with the goal of improving the speed and robustness of computational codes. However, truly predictive capabilities cannot be achieved only by improving computational techniques and algorithms, since the reliability of the predictions depends in a significant manner on the accuracy of the mathematical representation of the physical phenomena of interest. Therefore, there is a critical need for basic research in constitutive modeling, i.e., in developing the set of governing equations describing the strongly non-linear, history-dependent, large-scale, thermo-mechanical response of materials.

Among the routes in achieving weight reduction of aerospace platforms or improving robustness and reliability and increasing payload capacity are:

- Development of analytic yield criteria and plastic potentials accounting for the specific symmetries of each of the crystal classes.
- Implementation of multiscale analysis to reflect microstructure or nanostructure evolution at the continuum level in order to more accurately model material behavior.
- Development of improved engineering-scale plastic potentials for the description of the strong anisotropy in the mechanical behavior of textured metallic materials.
- Development of advanced models for prediction of damage evolution in ductile and brittle isotropic /anisotropic materials to improve material life prediction capability and minimize overdesign of load-bearing structures.
- Development of non-local constitutive models to better account for instabilities in materials and structures for both quasi-static and dynamic loadings.
- Development of multi-objective optimization methods for form analysis of the constitutive functionals.
- Development of clear and robust parameter identification strategies to ease application and transferability.

Development of such continuum level predictive constitutive models are key to achieving game-changing improvements in design for large-scale coupled multi-physics AF applications such as hypersonic vehicles design accounting for Plasticity-CFD-Thermal couplings; weapons design based on a better understanding of penetration events of a metallic systems into geologic/ultra-high strength cementitious materials; and detonation of energetic materials which involves strong coupling between the chemistry and the thermomechanical behavior of the constituent organic crystals.

Concerning discrete modeling, despite tremendous progress achieved in the past decade, the quality of the atomistic simulations is limited by the availability and quality of the interatomic potentials. There is a critical need for development of high-fidelity interatomic potentials, with emphasis on rational and reproducible development methods. A tremendous

challenge remains the establishment of protocols for transitioning data/knowledge between length/time scales.

Integration between modeling and experiments

Irrespective of the scale, an important topic of research remains the integration of experiments and simulations. For example, to take full advantage of the tremendous progress that has been made in experimental imaging techniques such as high-energy x-ray methods or digital image correlation, interpretation and analysis should be enriched through direct comparisons to models based on novel constitutive relations. Such models need to be developed in concert with recently developed and still developing synchrotron x-ray based experimental methods (where AF support continues to play a key role) that directly probe the nanoscale to macroscale responses being modeled.

Likewise, interpretation of mechanical data on the novel light-weight materials being developed can be done only through use of advanced constitutive models. However, a key challenge faced by modelers is that model development relies mainly on data for one-dimensional loadings (1-D), yet the models should be applicable for combined multi-axial loadings. There is a clear need for data/experiments to determine internal material length scales for various and multiaxial loading types/modes. To allow for verification/validation the following critical research directions have been identified:

- Development of standards to facilitate the collaboration between experimentalists and theoreticians/modelers for description of the behavior at lower scales.
- Development of comparison metrics that include concepts of statistical characterization.
- Development of standards and protocols for model calibration, verification, and validation at multiple length scale.

Computational Mechanics

Crucial to the design and control of future AF systems is the availability of robust and reliable computational capabilities to AF scientists and engineers. To this end, it is imperative that the development of new constitutive models should be tied to the development of appropriate computational frameworks for simulation of the thermo-mechanical behavior of any given system under complex loadings. Moreover, it should be ensured that such computational capabilities can be transitioned/integrated into both government and commercial codes through specific user element subroutines to characterize appropriately enhanced constitutive models. To fulfill this goal, it is imperative to develop robust algorithms with detailed documentation and associated benchmarks relevant to AF applications.

Concerning the multi-physics problems, which at present are considered as decoupled, it is important to define a common computational framework in which each set of constitutive equations describing the individual subsystems (e.g. penetrator and target subsystems in the case of penetration events, energetic material and casing materials, structural materials and

surrounding fluid environment, etc.) can be resolved accounting for the interactions between the respective constituents. This long-term overarching strategy is essential to filling the gap in virtual prototyping of AF systems.

The perennial challenge associated with quantification of uncertainties of the computational results should be addressed. Special emphasis should be put in quantification of the uncertainties that can originate from the form of the model itself (for example, the model doesn't include a critical physical mechanism), the identification of the model parameters, the quality of the experimental data, and the model used for interpretation of the data.

2. Advanced and Additive Manufacturing

Direct fabrication of functional end-use components and products, mainly represented by additive manufacturing (AM), is becoming the main trend of advanced manufacturing technology. The success with which 3D structural and functional parts consisting of multiple materials and functionally graded materials are being directly fabricated from various metallic, plastic, ceramic, electronic, composite, and/or biological materials has been increasing. However, AM will not replace conventional manufacturing methods for high-volume, low-complexity parts in the foreseeable future. Like many other manufacturing technologies, future AM technology will have more material options, better resolution, faster production, easier and more reliable operation, robust certification, cyber implementation, and lower costs.

For a manufacturing process to be adopted widely by the U.S. Air Force, it is critical to have repeatability and consistency of the manufactured parts for the entire build volume, between builds of each machine, and across different machines of the same make. AM technology must be able to guarantee consistent and reliable material properties in order to achieve wide adoption. Additionally, to broaden and develop new applications, novel AM processes, such as those for micro- and nano-engineering, must be investigated and developed. To achieve these goals, AM technology and its applications require significant further research and development in terms of design, materials, new processes and machines, process modeling, process control, and energy and sustainability applications. Towards this end, some major AM research issues and recommendations are summarized as follows:

- **Materials:** Development of new materials for AM processes, formation and mixing of materials in desired forms (powder, ink, slurry, etc.) and with desired properties (viscosity, solids loading, etc.), and understanding of the interaction between materials and processes.
- **Design:** Methods and tools for simultaneous multi-functional product design and AM process design, methods and tools for assessing lifecycle costs and impacts of AM fabricated components and products, and computer-aided design tools for representing complex geometries with repetitive features (e.g., cellular structures), multiple and gradient materials, and other variations.
- **Modeling:** Robust physics-based mathematical models of temperature, stress, and composition history for understanding transport phenomena in AM processes (e.g., melting and recrystallization), prediction of microstructures and fatigue properties resulting from extreme heating and cooling rates in AM processes, multi-scale modeling for more accurate prediction of complex process-structure-property relationships, modeling of layer-to-layer bonding between homogeneous and heterogeneous materials, and reduction of complex process models to lower-order models for the purpose of real-time control of AM processes.

- Sensing and control: Fast-response sensors for detecting defects and phase transformations, fast *in situ* temperature and cooling rate measurements for microstructure prediction, fast in-process optical measurement of geometric dimensions and surface quality of finished layers, in-process multi-band spectroscopic analysis of AM by-products, correlation between model predictions and online measurement of quality of AM-fabricated components, and integrated real-time sensing and closed-loop control of AM processes.
- Process innovation: Increase in AM material deposition rate by orders of magnitude for fabrication of large-scale parts and higher rate production needs, improved gas flow for cleaner surfaces during processing, novel bioprinting for simultaneous fulfillment of biophysical, biological, transport, anatomical, and manufacturability requirements, and digital fabrication with digital materials.
- Systems integration: Integration of interdisciplinary knowledge rooted in fundamental sciences for AM system development, and cyber-enabled and cloud-based sharing of AM and other manufacturing resources.

A tight coupling exists between material development, process development, process sensing and control, and qualification and certification of products fabricated by AM. As in the progression of many other emerging technologies, the greatest advancements will come at the boundaries of fundamental material science, physics, biology, lasers, electronics, optics, metrology (e.g., computed tomography, 3D scanning, etc.), and control.

The production costs, manufacturing time, and part defects must be reduced in order for AM to become an integral part of the U.S. Air Force manufacturing portfolio for production. The concept of “certify-as-you-build” is beginning to gain attention in industry; this robust and viable approach can cut costs, reduce waste, eliminate defects, and alleviate *in situ* process problems. It can be used in a variety of industrial sectors to revolutionize current manufacturing and inspection methods, overcome existing manufacturing challenges, and improve industry competitiveness. The outcome will be the elimination of waste and dramatically reduced lead times for the insertion of new products. This may result in a paradigm shift in the manufacturing industry and lead to the revival of manufacturing in the U.S., in particular, those for the U.S. Air Force. To provide the technological breakthroughs required to establish the new “certify-as-you-build” manufacturing paradigm, fundamental research (at TRL 1-2 levels) will be needed in order to achieve the *in situ* diagnostics and real-time control of AM processes. Integrating modeling, sensing, and process control will enable the “certify-as-you-build” vision for manufacturing processes.

3. Excited State Materials Performance

The Air Force needs robust and affordable technologies to support air-launched munitions to be used against hard, buried targets, ground targets and WMD facilities. Munitions must be designed to be more efficient and flexible, to enable multi-mission capabilities, with enhanced lethality and precision outcomes. The explosive loading output of a munition generates extreme states that in turn generate mechanical, thermal and chemical states that are beyond current resolved experimental measurement, and beyond the intended design limits for structures that interact with the explosive output. Simultaneous multiple effects can interact to alter material or system response. Here technical issues are identified as required to build such new near-term future capabilities. In particular, the focus is on the dynamic response of structures, and failure modes, and the design of reactive materials that include solid propellants and explosives. The challenges in this area can be summarized by the following questions: Can we design munitions with specific load functions to induce desired structural responses? Can we make predictive calculations for facilities subjected by such loads? Can design what we need by predictive calculation? Can we experimentally measure loads and response at extreme states?

In summary, the primary challenge of modernizing defense capabilities with emerging science lies in the fundamental understanding of the relationship between engineering design and function required for operation capabilities associated with the active engagement of potential adversaries. The objective of global technological superiority in the air, space and cyber space covers an impressive suite of science and engineering focus areas. At least one source defines engineering as the application of scientific and mathematical principle to provide a desired outcome.

Necessary Requirements for Predictive Modeling and Design

Addressing the previously-highlighted issues one must include a comprehensive understanding of the physics, chemistry, structural behavior, and measurement techniques of extreme states, as input to both high fidelity and fast-running numerical simulations. The material states are often characterized by pressures of up to many Mbars, with temperatures of up to 10,000 K, and the corresponding phenomena occur on a wide range of time scales (between nano-seconds to seconds), and in optically and electronically opaque environments. Currently, there is both a lack of physical capabilities to measure such phenomena and an absence of predictive theory for material and structural behaviors under such conditions. Modeling penetration and damage mechanisms for structures exposed to explosive system output, requires simulation of the states that are currently best estimated by extrapolation, and for which little experimentally validated theories exists. Further, one must consider a large variety of materials that are, or could be used for structural applications (i.e. pretty much any known material or material combination).

Challenges to Improved Explosive System Design

A major obstacle to improved explosive system design and controllable output is the lack of understanding of the chemistry of energy release. Since the initial energy releasing processes occur in dense, optically and electronically opaque, extreme states, direct measurement of thermodynamic and chemical states are largely not possible. One can measure macroscopic particle velocity at interfaces, but pressure, temperature and species concentration are inferred. The reaction zone of a high-performance explosive is often far less than 100 microns. Likewise, other explosives that are based on slurries or made from heterogeneous mixtures of explosive crystallites and binders and other additives, burn in reaction diffusion regions near the interfaces of chemically different mixture components. This is not modeled well if at all. Models that describe detonative processes, are quite under-constrained by experimental data, and thus model predictions are based on extrapolation of thermo-mechanical-chemical states into performance regimes where there is little data and hence no validated theoretical predictions. Improvements in measurement (i.e. diagnostic experiment) always precedes predictive (scientific theory) that later enables engineering calculation (i.e. modeling and simulation). Invention is required to find new ways to make measurements inside condensed phase materials (not optically accessible), and this includes consideration of x-ray, other electromagnetic and chemical signatures. To test theories, large parameter studies are needed. Performance studies depends on initial conditions, boundary conditions and loading conditions are subject to geometric complexity.

The following improvements in experimental measurement and capabilities could lead to significant improvement in predictive modeling and simulation:

- Development of "cheap" hard x-rays for imaging in condensed phase media such as explosive, i.e. "X-Rays on a Chip." At the present time, short duration shock experiments require a synchrotron (Advanced Photon Source). Therefore, the invention of a hard x-ray source that was easily focused, would obviate the costs of a large accelerator source. Cheap hard x-rays might be able to be generated with a small explosive flux compression generator that drives energy into a target sample that generates the necessary x-rays.
- Development of small scale, low-cost experiments for determination of the shock equation of state for materials at extreme states.
- Development of high-fidelity (i.e. accurate physics) for reactive molecular dynamics for determination of: Phase transformation, Reactive intermediates, Reaction kinetics schemes and rates, and Multicomponent mass and thermal diffusion.

While the development of operational capability based on the acknowledged emerging technologies might be considered an engineering task linking capabilities to technology; the first task is defining the underlying science and mathematics, together with their strength and weaknesses in terms of technologies listed together with their stated mission. The technological superiority objective requires the ability to reach into the heart of the technology and unravel the science that defines the limits of capability, cross-cutting advancements in physics, chemistry,

mathematics, and translation into engineering application to the functional attributes of new capability. There are numerous examples of new technology reaching out for advanced engineering to provide affordable edge cutting [products, or profit] such as digital television and the spread of cellular communication. Performance as a measure of capability is quantified in the evolution of capabilities at the sub-system level in complex infrastructure characterized as a system of systems.

As an aid in assessing performance and impacting the affordability of prototyping; modeling and simulation (M&S) tools are employed assessing the functional attributes based on design information. Several decades of development of M&S tools have been jointly developed under the umbrella of a memorandum of agreement between the Department of Defense and the Department of Energy; the management charge of the nation's nuclear weapons complex. Central in the transfer of capabilities has been the developed codes that are functional on state-of-the-art super computing networks and extensive continuum base modeling of a wide variety of materials comprising what may be characterized as structural, architectural, geological and energetic materials ranging from explosives to propellants. The validity the assessment is founded upon the fidelity and resolution of the M&S tool used in the assessment. The input of the "design" information is based on the complete suite of engineering specification documented on materials and processes used to manufacture.

The design information is the input to the M&S analysis as a means of quantifying capability in terms of systems, sub-systems, component, and sub-component type analysis, to evaluate the functional attributes of design. The analysis also aids in disclosing the design elements that limit the utility and ultimately the capability in the context the overall system function.

In reviewing the list of emerging technologies cross cutting areas of science, mathematics and physics and chemistry, are identified. Experience and experimentation identifies basic science problems associated with the distinction between fidelity and resolution. We identify the engineering product the design development as being a complete suite of specifications on materials and processes used to prototype the technology for the desired application together with a rather extensive list of operational specifications like affordability, safety, etc. The environment on the battle-field together with the proper function of materials in a weaponization context place material under extreme loads. Current specifications on materials together with emerging manufacturing capabilities increasingly define small scale atoms and molecules) that determine the physical capability of the end product in these conditions.

A careful review the Military Standard specification on the ingredients list used in the formulation of high explosives and propellants is exemplary of the element of engineering design of payloads used in the fight. There is an indication that the improved science supporting higher fidelity in the physics and chemistry of the M&S approaches used in evaluating engineering designs not only in the class of energetic materials but structural materials as well. Part of the goal is always to do more with less which means miniaturization of the overall munition system. The value of advanced materials impacts both offense and defensive need on the battle-field.

In conclusion, a major shift in the physical scale of modeling and simulation will be required to address the physical scale of engineering design based on specifications using emerging manufacturing technology and miniaturization. Historically, computer-aided design was pursued with a passion, driving the machines to the massively parallel networks of today while focusing on the numerical methods to make large codes such as ALE3D, CTH and EPIC run fast with DOE investing in programs such as ASCII and PSAD, without the focus on advancing the science that determines the function of the end product. In contrast, the fact that the molecules and the atoms used in the design are in fact quantum entities has been overlooked, which in turn impacts the function of the larger scale end product. The fundamental foundations of engineering are the quantitative assessment of the relationship between design and function. The community must continue to address the challenges of so called nano-materials and micro-mechanical systems supporting miniaturization of various commercial and military hardware.

4. Energetics, Fuels and Propulsion

This section discusses the research that is necessary to develop next-generation propulsion and power systems using fuel sources, as divided into two primary sub-sections. The first section discusses combustion technologies for propulsion and power generation. The second section discusses alternative uses of fuel such as fuel cells for small portable power sources.

Combustion-Based Power Sources

The objective and importance of combustion-based energy sources is to develop next-generation combustion technologies that operate with higher performance, increased efficiency, lower operating cost, and reduced emissions. To achieve these objectives, next-generation combustors are being driven to higher pressures, higher levels of dilution, leaner mixtures of fuel and oxidizer, and utilizing non-traditional modes of combustion such as detonations. Combustion theory, however, has traditionally been focused on burning fuel and oxidizer mixtures at lower pressures and typically flaming modes of combustion. The operating conditions of next-generation combustors are stretching the limits of our knowledge of chemically reactive flow and may even require rethinking some of the underpinning assumptions of modern combustion theory. As a result, there is a large gap in knowledge that needs to be filled before these next-generation combustion technologies can be developed.

Recommendations for research that will enable next-generation propulsion, rocket, and power generation combustion technologies for 2030 and beyond are:

- Significant research effort is needed to develop chemical reaction, transport, and large-eddy simulation models at supercritical conditions.
- Novel experiments are needed to measure the bulk combustion properties of fuel and oxidizer mixtures, such as the laminar flame speed, at high-pressure conditions.
- Develop experimental and numerical simulation techniques to allow the fundamental study of liquid-fuel combustion high-Mach number and shock-laden flows typical of hypersonic propulsion devices and detonation engines.
- Develop novel experimental approaches to calibrate chemical reaction models that eliminate measurement uncertainties associated with non-uniform ignition of the fuel and oxidizer mixture.
- Explore next-generation combustion technologies such as pressure-gain, MILD, and HiTAC combustion to allow for higher efficiency, reduced pollutant formation, and potentially fuel-flexible combustors. Significant research in developing chemical reaction models, optically accessible experiments, and advanced diagnostics is necessary to fully realize these concepts.
- Continue research in traditional, but unsolved areas of combustion. For example, combustor instabilities related to combustion-acoustic coupling remains an enormous problem when designing combustion systems.

These points are discussed below.

Combustion at High-Pressure Conditions

The operating pressure of combustors for propulsion and power generation devices is increasing in an effort to reduce system size and increase efficiency. For example, supercritical CO₂ cycles




Figure 1: Laminar flame speed of mixtures of methane and oxidizer mixtures at 60 bar as function of equivalence ratio. The oxidizer is a mixture of oxygen and helium at molar ratios of 0.15:0.85. The lines represent the laminar flame speed computed using different chemical reaction models and the circles are experimental measurements. (Strakey, 2018).

for power generation are expected to operate at conditions in excess of 100 bar (Strakey, 2018). These high-pressure conditions are stressing the limitations of chemical reaction and transport models used to design and understand combustion systems. Many chemical reaction models were not designed or calibrated for such high-pressure conditions. It was shown that the computed laminar flame speed at high pressure has substantial variation when different chemical reaction models at high pressures are used (Strakey, 2018). Figure 1 shows the laminar flame speed for mixtures of methane and oxidizer at 60 bar computed using several different chemical reaction models. It is clear from this Fig. 1 that there is a large amount of disagreement between different chemical reaction models at high pressure. More importantly, none of chemical reaction models agree with the experimental data over the full range of equivalence ratios.

High-pressure combustion systems, such as liquid rocket and diesel engines, operate at even higher pressures where the liquid is in a supercritical state. At these conditions it is well known the equation of state and transport properties need significant correction (Yang, 2000). Numerical simulations show that the mixing of fuel and oxidizer at these supercritical conditions can be highly sensitive to the choice of molecular transport model. In addition, many traditional

large eddy simulation closure models are inadequate at these high-pressure conditions, which is likely due, in part, to nonlinear equations of state (Borghesi, 2015).

There has been very little discussion in the literature regarding the chemical reaction models at supercritical conditions. From a fundamental perspective, the binary collision assumption, which is implicit in most chemical reaction models, is, at the very least, questionable at high-pressure and supercritical conditions. As a result, the applicability of traditional Arrhenius chemical reaction models needs to be explored and, if needed, corrected or replaced by a more accurate models.

Developing combustion technologies at these high pressures requires the development of accurate chemical reaction models for high-pressure conditions, the development of turbulence closure models compatible with highly nonlinear equations of state. Additionally, novel experiments that focus on measuring bulk combustion properties, such as the laminar flame speed, are needed at high pressure to verify numerical simulations.

Novel Combustion Concepts

There are too many novel combustor concepts for next-generation combustion systems to highlight in this report. Instead the focus is on two concepts which show promise for different applications. These include Pressure Gain Combustion and Moderate or Intense Low-oxygen Dilution (MILD) combustion.

Pressure gain combustion (or detonation engines) is a promising technology that offers increased thermodynamic efficiency over traditional flame-based combustors (Wolanski, 2013). The combustion process, usually a detonation, directly increases pressure in these devices. This in contrast to tradition combustors where pressure is increased from a combination of mechanical compression and the conversion of dynamic to static pressure.

The extreme conditions inherent to pressure-gain combustors make fundamental studies using either experimental or numerical simulation techniques difficult. Powerful shocks propagate throughout these combustions, the pressures high, the time scales small, and these devices are optically inaccessible. Improvements in experimental measurement techniques and novel experimental test rigs are necessary to increase fundamental understanding of detonation-based combustion technologies. Linear detonation engine rigs (Burr, 2017) are a good start. Nevertheless, it is necessary to develop improved optical diagnostic measurement techniques in the highly compressible and shock-laden flows characteristic of a pressure-gain combustor.

Future Air Force missions may be affected by wildly varying local fuel sources. For example, a wide variety of bio and synthetic fuels are under development and the local type and quality of fuel may vary in the future. It is advantageous that the Air Force develop combustion technologies that can utilize these local fuel sources with lower pollutant formation. There are several combustion technologies that been proposed including high temperature air combustion (HiTAC) and MILD combustion (Khalil, 2013). However, much research is needed to advance these combustion technologies to the performance requirements of military air craft.

Combustion of liquids at high-speed and shock-laden flows

Many practical combustors for high-speed and hypersonic propulsion systems, such as scramjets and pressure-gain combustors, will likely use liquid hydrocarbon fuels. Liquid hydrocarbons have an unrivaled combination of energy density, ease of storage, and transportability. Development of liquid-fueled pressure-gain combustors or hypersonic propulsion devices requires understanding of how fuel sprays and droplets atomize, vaporize, and burn in high-Mach number and shock-laden flows. However, almost all of the academic research for these devices centers around gaseous fuels such as hydrogen or ethylene.

Understanding the combustion of liquids in these advanced combustion devices will require a combination of detailed numerical simulations and clever experiments. Advances in both numerical simulation techniques and experimental diagnostics are required to achieve this understanding. Experimental diagnostic techniques need to be developed for these challenging and often optically inaccessible environments. Numerical simulations of these flows are also challenging and very few simulation algorithms have been developed to study the combustion of liquid fuels in shock-laden flows (Houim, 2013).

Practical application of liquid-fuel combustion in high-speed and shock-laden flows is expected to be restricted. This is especially true for hypersonic propulsion devices such as scramjets. However, universities have much to offer in producing the fundamental and basic knowledge that will undoubtedly serve as the building blocks needed for the Air Force and its contractors to develop these devices. The development of open geometries or studying canonical problems, such as the combustion of a liquid fuel spray in a supersonic cross flow, will allow universities to contribute research and train future DoD scientists and engineers in this area.

Experimental Development of Chemical Reaction Models

Developing next-generation combustion technologies require knowledge of chemical reaction mechanisms at high-pressure, heavily diluted, low fuel concentration, or highly compressible flow conditions. Most of the understanding of chemically reactive flows has been developed for conditions typical of traditional combustion technologies. Current chemical reaction models are simply not applicable to many of the combustion technologies highlighted above.

It has been shown that detailed chemical reaction models perform poorly and do not even qualitatively agree with experiments when used to simulate hydrogen-air detonations at realistic pressures (Taylor, 2013). One reason for this poor performance is the introduction of thermodynamic nonequilibrium in these strongly-shocked flows. Figure 2 shows another reason for this poor performance. There is simply very little overlap of the experimental conditions used to calibrate chemical reaction models at the pressures and temperatures relevant to a hydrogen-air detonation. Thus, simulations of next-generation combustion technologies must extrapolate these highly nonlinear chemical reaction models far beyond their calibration regime. This issue is exacerbated for more complex hydrocarbon fuels that are of practical interest to the Air Force.




Figure 2: Results from a numerical simulation of a hydrogen-air detonation. The red dots represent a scatter plot of pressure and temperature of unburned material from the simulation. The black and green dots represent the points in pressure-temperature space where experiments were to calibrate chemical reaction models. (Oran, 2015).

Figure 2 and the discussion above highlight the need for the development chemical reaction models at the high-pressure, highly dilute, low equivalence ratio, and shock-laden flows that may be typical of advanced combustion systems. However, experiments at these conditions are very difficult and commonly-held assumptions may no longer be accurate. For example, it is generally assumed that the reactants ignite uniformly in experiments aimed at calibrating chemical reaction models. However, recent experiments show that this is not always the case (Pryor, 2017). Figure 3 shows that the fuel and oxidizer mixture can ignite in small localized regions and then propagate as flames or detonations. This issue can lead to significant challenges and errors when these experiments are used to calibrate chemical reaction models. Clearly, more research is needed to develop novel experiments that eliminate this problem.




Figure 3: High-speed photographic images of the non-uniform ignition of reactants in a shock-tube apparatus used to calibrate chemical reaction models (Pryor, 2017). The reactants are a stoichiometric mixture of methane and oxygen that is heavily diluted by carbon dioxide and argon. The mixture ideally ignites uniformly; however, high-speed images show that the ignition is highly non-uniform.

Alternative Power Sources

Future Air Force mission goals will no doubt expand toward longer duration and range, in addition to increased reliability, lower cost, and reduced dependence on materials sourced from politically unstable regions. To achieve these goals, the most sensitive engineering metrics are improved energy density and power density. For example, small unmanned aerial vehicles (UAVs, or drones) are currently battery powered due to the poor power and energy density of competing energy conversion technologies in the few hundred Watt size class and below. As such, small aircraft have been designed, developed, and implemented based on batteries. Batteries have the advantage of high power density, which allows for necessary power output required for take-off and acceleration. Batteries, however, suffer from low energy density, resulting in insufficient flight duration without a dramatic increase in airframe size. Given that battery technology is fairly mature, this limitation of battery technology is expected to improve only incrementally, and thus superior small-scale energy conversion technologies are necessary to meet future goals.


Most of the future mission goals for UAVs, as one example, could be achieved if an ultra-high energy density power source were available. An airframe designed around such a power source could remain compact (man portable) while allowing continuous flight for many hours or days. This scenario not only would provide a significant operational and tactical advantage, but the smaller aircraft size would reduce cost and support requirements.

Any portable power application for which size or weight is a design driver, combined with longer mission duration, would similarly benefit from the development of ultra-high energy density small power sources. These applications could include any type of unmanned vehicle, mobile robots operating away from the power grid, remotely-based electronics for sensing, or portable battery chargers.

Fuel Cell Technology

Fuel cells operate from the same electrochemical process as batteries, but have crucial advantages for long duration missions. Fuel cells utilize externally-stored fuel (hydrogen, methanol, etc.) in a power-producing module to produce electricity while battery energy storage is internal, and relies on strategic materials such as lithium and cobalt. As a result, fuel cells have a low power density but comparatively high energy density as the mission duration increases.

Lithium-ion batteries have an energy density of approximately 200 W-hr/kg. Fuel cells, theoretically, can achieve energy density many times that value. For example, Direct Methanol Fuel Cell (DMFC) systems developed by University of Florida (UF) and University of North Florida (UNF) have achieved an energy density of close to 800 W-hr/kg during testing by the U.S Army CERDEC (Thampan, 2014), shown in Fig. 3. The expected maximum efficiency of that system when mature is 1200-1500 W-hr/kg, which could result in a significant increase in mission duration for the same UAV size.



UF/UNF Direct Methanol Fuel Cell System. Energy density 792 W-hr/kg for a 72 hour test by U.S. Army CERDEC.

Advanced Concepts

There are many advanced fuel cell concepts that have potential to significantly improve energy density. The DOD-funded system referenced above utilizes a passive water recovery concept that minimizes the weight and volume associated with the power conversion module, thus achieving high energy density. Other concepts include 1) utilization of high energy density fuels, such as hydrides, 2) miniaturization of solid oxide fuel cell systems, and 3) lithium-air batteries.

Hybridization

As mentioned above, fuel cells typically have the potential for high energy density but suffer from low power density. For this reason, among others, hybridization is often used, so that the fuel cell power supply includes a small battery storage subsystem, analogous to a hybrid automobile. Short duration high power demand is met by the battery, allowing the fuel cell to be

sized for the average power requirement, not the peak. For example, the DOD-funded system described above includes Li^+ batteries which occupy less than 5% of the total package volume.

Future Needs

High Energy Density Power Sources - For long duration missions, a five-fold increase in energy density would allow true breakthroughs in operational capability. Hence a target energy density in the 1000 W-hr/kg range would seem to be an appropriate and achievable goal. These advanced power sources also will be required to benefit (or at least not degrade) the remaining future mission goals, such as strategic materials requirements, reliability, and cost.

Overall System Design - A key research topic is to investigate optimizing the integration of fuel cell technology into small air vehicles, rather than simply attempting to integrate fuel cells into existing designs based on battery technology. Proper integration will result in maximizing the advantages of fuel cells, particularly for the desired long-duration missions. For example, minimizing the energy consumption during level flight through choice of components such as electric motor benefit fuel cells. The system design for a long duration UAV must recognize the relatively smaller time and energy consumption due to takeoff and acceleration, relative to the energy used in level flight. Much as commercial aircraft are optimized based on the two flight regimes, the UAV design must also be optimized, and that optimal design will likely differ significantly from the current battery powered designs.

5. Hypersonics and High-speed Effects

The hypersonic regime is characterized by overlapping entropy and boundary layers. For flight in the earth atmosphere this corresponds to Mach 5 and above. Vehicles flying in this regime must manage large amounts of energy and a broad range of operational condition. Together, these flight characteristics impose severe operating environments for flight surfaces, propulsion systems, systems and control. Major current concerns for design and development were discussed during the S&T Summit as summarized here. The discussions focused exclusively on the technical issues facing the development of hypersonic flight and intentionally excluded references or suggestions to mission requirements. The relevant consensus recommendations are listed further below.

Findings and Recommendations

The group identified several areas that continue to require technical development for implementation in the design of the hypersonic vehicles:

- Development of Materials that will extend the operational envelope of hypersonic systems
- Propulsion systems integration and operability
- Analyses capabilities:
 - a. Theoretical analyses
 - b. Experimental testing
 - c. System integration.


The consensus on these requirements are discussed here for each of the three topics in more detail.

Materials

Large components on the hypersonic vehicle external and internal surfaces will experience significantly elevated temperatures. Figure 4 is an example of expected temperatures at Mach 7 and 10 along with material selections for the experimental X-43 vehicles. It should be noted that the X-43 vehicles flew for only several seconds. Clearly, new materials and their management will be needed to survive sustained flight at hypersonic conditions thermal loads. Thus, the consensus of the group was that the following requirements exist:

- Ultra-High Temperature materials that will be capable to operate continuously at temperatures in excess of 4000°F.
- These materials will have to resistant to thermal shocks that will result from formation and displacement of strong shock waves.
- Certain region of the vehicle will have to be designed for thermal management. For example, leading edge heating will be severe and some form of heat rejection will become necessary either in the form of active cooling or via the use of ablative materials.

- Along with the development of new materials manufacturing techniques suitable for these compounds will have to be developed.



Propulsion System and Integration

Given the broad range of flight regime no single known propulsion cycle will be capable to propel the hypersonic vehicle throughout the entire operational envelope. Aside from the costly option of staging, combined cycles will need to be developed. Figure 5 shows an example of turbine-based combined cycle and a set of technological issues to be resolved. Other combined cycles that have been proposed and deserve further study include rocket-based combined cycles and inclusion of other options such as constant-volume combustion for the gas-turbine component of the combined cycle. Major technological issues to be solved include:

- Engine flow path-vehicle integration: since the entire lower surface of the hypersonic will be occupied by the propulsion system, the development of the engine cannot be uncoupled from the vehicle design and operability over the entire envelope including changes in center of gravity, forces and moments resulting from transient operation and the maneuvers required through the atmosphere. There is a need to further expand the current scientific and technological understanding in the specific following areas: Aerodynamic boundary-layer transitions; low pressure/moderate temperature combustion; and unstart control, etc.

- Combined cycle performance and operability, i.e., transition from turbine-only to ram-scramjet operation has to be ensured with continuous modulation of thrust along the flight path.
- Inlet/isolator/combustor stability must ensure unidirectional airflow, i.e., eliminate the danger of inlet unstart or engine flameout as transition between power modes occurs and the ram/scramjet combustion chamber experiences transients.
- The design has to ensure the stability of the operation described above while ensuring the operability required by the mission.




Fig. 5. A proposed combined turbine-based cycle example and technological issues to be solved. (From *Defense Industry Daily*, Jul 05, 2010).

- Combustor operational envelope in ram/scramjet mode will be required to ensure stable operation from low to high speed and produce thrust efficiently. Structurally, the combustion chamber will be subject of high thermal loads that need advanced active cooling method and technology demonstration.
- Flameholding is a topic for which predictive, knowledge-based, technological methods remains elusive. Although it is understood by the community that formulation of flameholding must be in terms of the D amkohler number, a relative comparison of chemical reactions and residence times, predictive tools under the entire operational envelope do not exist; currently, every design is based on direct analysis.

Experimental Testing

Ground testing remains a difficult problem given the large power installations necessary to simulate flight conditions. As the hypersonic designs have evolved several requirements appeared necessary:

- “*Test like we fly*”- indicates the need to develop clean, continuous, non-vitiated facilities for the testing, in particular for propulsion systems evaluation. To achieve stagnation temperatures/enthalpies experienced in hypersonic flight vitiation resulting from the air heating method used in many current ground facilities is likely to change the air composition to a degree that invalidates the extrapolation of the results for flight. Short term facilities, such as shock tubes, that can reproduce very large flight Mach numbers fail to reproduce the thermal regime under long duration operation.
- Flight testing will remain a necessary component in the development of hypersonic vehicles. New technologies are needed to perform these flights in a routine, affordable manner. Some of the options envisioned by the committee include:
 - Sounding rockets with staging ability,
 - Configurable for mission requirement.
- High-temperature instrumentation will be required both for testing and for implementation in the actual flight object. These instruments will require high-frequency – given the high-speed flight regime – and new packaging concepts, to respond to the demands of compactness of design and lightweight systems components.
- Communication and telemetry has to be developed along with the instrumentation both to monitor/control the flight and to survive under possible enemy adverse interference.

Operational-Specific Technologies

As other devices, the hypersonic vehicle will rely on additional components needed to sustain the airframe/propulsion system. These auxiliary systems will need to be developed specifically for the conditions of high-speed, severe thermal loads and broad operational conditions. Among the technologies identified for future development the committee recommended the following:

- Power generation for A/C components.
- System integration/optimization/controls for Maneuverability at high-speed.
- Component interface definition at the aerodynamic interface plane, but also development of modular airframe concepts capable of staging as future missions will evolve.

6. Cyber Systems

The ability to execute autonomous operations or to achieve multiplied force projection through networked systems depends on the integrity of the information being shared and on the hardware that it is being implemented. Researchers from industry and the Florida Institute for Cybersecurity hosted breakout sessions during the summit and pre/post summit meetings to discuss cyber effects that could compromise the integrity of military systems. This section summarizes these cybersecurity topics in sections including: Novel Cyber-Resilient Computing Architectures, Standardized Technology Requirements for Trusted Authenticity in Supply Chain Assurance, Anti-Tamper Technology Protection of Critical Assets from Physical Attack, and Cyber Safety in Networked Communications.

Novel Cyber-Resilient Computing Architectures

For the past seventy years, the foundation of modern computers has been the von Neumann architecture. However, recent vulnerabilities in pervasively deployed processors from Intel (Meltdown and Spectre) and AMD (Masterkey, Ryzenfall, Fallout, and Chimera) demonstrate that it may be time to rethink computer architecture and organization, specifically from a security-oriented perspective. Today's computing systems are complex, having grown from single-chip processors to multiple cores in the last two decades. Modern processors are also less self-contained, and process multiple tasks using shared resources (cores, memories, caches, buses, etc.). The aforementioned attacks exploit inter-process or intra-process data leaks to extract secrets from processes running on the same physical machine without ever breaching the victim process. In the age of the Internet and cloud computing environments, this is particularly dangerous.

Although patches/fixes have served the software community well, they have failed to address the above issues. Spectre is more or less unpatchable because it is a class of attack that exploits branch prediction and speculative execution, which are baked into the hardware. Patches from Intel and Microsoft have also drastically compromised performance or caused other unintended consequences. Another major obstacle is that of "shared responsibility". In the case of cloud computing, even the best and fastest patches are only effective, if all stakeholders (cloud providers, OS vendors, hardware vendors, etc.) do their part - exposure shall remain if there is even one single weak link. Hence, these recent security/privacy issues can only truly be resolvable by new hardware-based solutions.

Going one step further, it is even more likely that similar issues will persist until we completely depart from traditional computer architectures and the economic forces that have driven them – simplicity, usability, and performance. Instead, novel computer architectures need to be investigated with security and correctness in mind from both the top down as well as the bottom up. The following problems need to be considered in future research:

- Coming from the top, security properties should be defined and integrated with other fundamental design principles. All types of architectures should be explored and then analyzed based on these properties. Examples include near data processing where movement of data is minimized as well as unconventional computing methodologies such as chaos, quantum computing, neuromorphic computing, stochastic computing, and

multiparty computation (e.g., homomorphic computing, secret sharing, and garbled computing).

- Consideration from the bottom up is just as important. For instance, recently developed Trusted Execution Environments (TEEs) guarantee code and data loaded inside to be protected with respect to confidentiality and integrity but only under certain assumptions. TEEs and other computing architectures remain vulnerable to a variety of physical attacks including side channel analysis, fault injection, and micro-probing. Therefore, the underlying implementation is just as important as the architecture itself. Ideally, emerging devices for nonvolatile storage and logic based on metal oxides, spintronics, phase change materials, graphene etc. will be explored to support and reinforce the security principles of new computer architectures rather than to only provide more capacity and/or performance.

TEEs enforce strong isolation guarantees against malicious insider attacks from untrusted system software, such as the operating system (OS) and/or hypervisor. However, almost all existing TEEs suffer from architectural side-channel attacks. The recently introduced architectural side-channel vulnerabilities, e.g., Spectre, Meltdown, and 4K-aliasing, have raised threats to the integrity and trustworthiness of TEEs. Further research is needed to explore more solutions, that can counter the architectural side-channels and to help secure the TEEs.

Security assurance at the platform level entails systematic exploration of vulnerabilities that cross-cut hardware, firmware, and software components. Unfortunately, this process today is purely driven by human insight to identify potential attacks based on past experiences, related intrusions, and even vulnerability research. As people get to an environment involving billions of computing devices (of which a significant fraction is potentially malicious) and rely on electronic security for more and more critical functions (e.g., automotive systems, military applications, medical devices, food authentication, etc.), further research is critically needed to develop more systematic vulnerability detection technologies. Particularly, innovative solutions are needed to predict security attacks at the platform level as early in the system life-cycle as possible.

System-on-Chip (SoC) security architectures targeted towards diverse applications including Internet of Things (IoT) and automotive systems enforce two critical design requirements: in-field configurability and low overhead. To address these constraints, further research is needed to develop a novel, flexible, and adaptable SoC security architecture that efficiently implements diverse security policies. The architecture and associated CAD flow should enable “hardware patching” i.e. a hardware security policy engine that can be seamlessly and securely upgraded in field to address unanticipated attacks or new security requirements. Innovative solutions are needed to fulfill the following requirements: (i) a centralized Reconfigurable Security Policy Engine (RSPE), (ii) smart security wrappers, and (iii) Design-for-Debug (DfD) infrastructure interface as the building blocks of the architecture. (iv) a systematic approach to represent and synthesize diverse security policies. Future research should also consider the evaluation of architecture on industrial SoC models with test chip fabrication and explore its potentials in untrusted 3PIP verification.

Standardized Technology Requirements for Trusted Authenticity in Supply Chain Assurance

Semiconductor intellectual property is undoubtedly one of the most valuable assets for any electronic design company, requiring enormous amounts of time, monetary investment, research, and development. However, such IPs can be compromised at various steps in the design process, most notably during fabrication at an untrusted foundry, and in the supply chain by reverse engineering attacks. Logic encryption has emerged as a promising solution that enables an IP owner to ‘lock’ their IP so that unauthorized parties are unable to engage in piracy or theft. However, such techniques are vulnerable to a wide array of attacks that either circumvent the locking mechanism or compromise the key required to unlock the IP or IC. In CHES ’17, we showed that such locking techniques, despite being secured by key protection techniques, are vulnerable to functional circumvention or bypass attacks. In light of this and numerous concurrently proposed attacks, further research is needed to: (i) evaluate logic locking against different attacks; (ii) evaluating trade-offs in overhead vs. security, (iii) extending combinational locking techniques to sequential circuits, and vice versa.

Deprocessing of ICs historically employs a variety of mechanical and chemical process tools in combination with one or more imaging modalities to reconstruct the IC architecture. An extensible programmatic workflow can take advantage of evolving technologies in 2D/3D imaging, distributed instrument control, image processing, as well as automated mechanical/chemical deprocessing technology. Ultra-thinning the silicon substrate in the packaged device within 1-2um of the IC device significantly reduces the amount of time required for deprocessing. The computer aided backside ultra-thinning approach not only improves the success rate, as compared to manual techniques, it also allows the dense lower layers with smallest feature size to be imaged via high resolution SEM first, while the sample layers are the most uniform. Backside deprocessing has the additional advantage that it can be possible to access the device while keeping it alive for in-situ electrical testing. To realize the above-mentioned solutions, innovative solutions are needed to enhance the deprocessing workflow with intelligent automation by bridging FIB-SEM instrument control and near real-time data analysis to establish a computationally guided microscopy suite.

In the past decade, electronics counterfeiting has reached a new level of sophistication. With improved replication techniques and the growing reliance on electronic systems worldwide, the possibility of important systems becoming compromised increases every day. In 2011, it was revealed that the Navy’s submarine hunter aircraft had compromised ice detection systems. Additionally, in 2014, the Navy’s nuclear submarines were discovered to use counterfeit electronic parts. Because these parts can allow backdoor access, it becomes prudent to vet the parts being used in important systems. Standard procedures involve manual examination of parts for defects by an SME, which can be expensive, time-consuming and destructive.

With the goal of automating the entire process, surface defects are needed to be detected with image processing techniques. To generate general solutions, further research is needed to: (i) detect scratches, texture differences on the surface of the chip, displacements in the identification marking locations, and color variations; (ii) develop an entirely automated process that can analyze whole trays of chips at once.

Globalization of design, fabrication, and assembly of electronic devices and systems and the rise of internet of things has raised serious concerns about the security and trustworthiness of integrated circuits (ICs) and electronic systems. In this project, we will develop a holistic end-to-end solution toward security of devices, systems, and firmware. There are many attacks on the microelectronic devices from the first step of the design process to end of life. Example attacks include IP piracy, IP security and trust, CAD tools induced security vulnerability, SoC security, hardware Trojan insertion, overproduction, shipping defective and out-of-spec parts, counterfeiting, reverse engineering, side channel attacks, and recycling. To mitigate all these potential vulnerabilities, problems need to be solved in the following two perspectives: (i) backward trust, where a SOC design house must ensure all IPs are secure and trustworthy. (ii) forward trust, where all entities engaging with the SOC design house are considered untrusted and must be prevented from performing piracy and tampering. In this project, innovative solutions will be developed to protect the design against side-channel attacks, fault injection, etc.

Semiconductor design houses are increasingly becoming dependent on third party vendors to procure IPs and meet time-to-market constraints. However, these third-party IPs cannot be trusted as hardware Trojans can be maliciously inserted into them by untrusted vendors. While different approaches have been proposed to detect Trojans in third party IPs, their limitations have not been extensively studied. Further research is needed to develop Trojan detection framework based on information flow security (IFS) verification in system-on-chip (SoC) designs. In particular, solutions are needed to detect the violation of IFS policies caused by Trojans without the need for white-box knowledge of the IP, as well as ensure the security and trustworthiness of millions of IPs used in the market.

Anti-Tamper Technology Protection of Critical Assets from Physical Attack

Re-routing internal paths and adding components on printed circuit boards (PCBs) makes it possible to bypass the copyright verification, break the carrier restriction, and run 3rd-party systems. The primary method of achieving these goals is to install a modification chip (mod-chip) into the system. A mod-chip is a small electronic device used to alter or disable artificial restrictions of computers or entertainment devices (e.g., videogame consoles, Blu-ray players, etc.). Physical countermeasures have been proposed to address this attack. For example, PCBs can be protected in a hard steel case with tamper evident switches that activate to suspend the system if the case is opened. Vibration sensors can be used to monitor abrupt movements. However, these solutions present some drawbacks. To eliminate tampering attacks, problems to investigate are: (i) how to capture unclonable signatures in the PCB traces based on process variation; (ii) how to build a general PCB attestation framework to monitor the changes introduced by tampering of critical board-level interconnections.

Physical attacks on electronics have grown significantly recently and is becoming a major threat to the chip designers, original equipment manufacturers as well as end users. The complex long life of the emergent internet of things (IoT) devices coupled with their diverse applications is making them increasingly vulnerable to various forms of physical attacks. Large number of companies and start-ups are now working in this area across the globe and offer training and service to the public. High-end failure analysis tools, which used to be very expensive and accessible to only big organizations are now available for a very cheaper price. Private parties

can now buy such tools and rent them to the public with a very low rate, which makes these attacks even more viable.

In addition to physical attacks, physical inspection of electronics as well as physical fingerprinting based on analog parameters are rapidly becoming a tool for unique countermeasures against various attack modalities. Therefore, innovative research is needed to analyze and find solutions for, side-channel attacks, reverse engineering, physical tampering, micro/nano probing, and physical access at chip, PCB and system level.

Side-channel information is collected from the physical implementation of a circuits, and side-channel information can be used to break into the system, obtain secret information, etc. Typically, timing information, power consumption, temperature tracking, electromagnetic leaks or even sound scan can provide an extra source of information of the circuit. While on other aspects, side-channel information is utilized by Trojans to leak specific data out of the circuit or is used as a source of fault attack. Further research is needed on EM side-channel information detection and vulnerability analysis: (i) establish a more precise multi parameter side channel collection system; (ii) analysis the relationship between different side channels; (iii) utilize the side-channel methods for integrated circuit security analysis; and (iv) find design flaws through side-channel analysis.

Cyber Safety in Networked Communications

There is an increased need of novel research on threats and challenges in the context of cyber safety, especially in modern online social networks, where the range of malicious activities perpetrated by malevolent actors is regrettably wide. These include spreading malware and spam, controlling and operating fake/ compromised accounts, artificially manipulating the reputation of accounts and pages, spreading false information as well as terrorist propaganda. The reasons for the success of such attacks are manifold. The users of social networking services tend to extend their trust into the services and profiles of their acquaintances to unknown users and other third parties, despite the service providers' attempt of keeping their audiences identifiable and accountable, creating a fake profile, also in another person's name, is very simple. Even partially or fully taking over a profile is comparatively easy, and comes with the benefit of the trust this profile has accrued over time, as many credentials are easy to acquire.

Further, even seemingly innocuous issues such as the design and presentation of user interfaces can result in implications for cybersafety. The failure to understand the interfaces and ramifications of certain online actions can lead to extensive over-sharing. Even the limited information of partial profiles may be sufficient for abuse by inference on specific features only. This is especially worrisome for new or younger users of a system that might unknowingly expose information or have unwanted interactions simply due to not fully understanding the platform they are using.

Unfortunately, past research in Cyber Safety has only looked at the various sub-problems in isolation, almost exclusively relying on algorithms aimed at detecting malicious accounts that act similarly, or analyzing specific lingual patterns. This ultimately yields a cat-and-mouse game, mostly played on economic grounds, whereby social network operators attempt to make it more

and more costly for fraudsters to evade detection, which unfortunately tends to discard to measure and address the impact of safety threats from the point of view of regular individuals.

This prompts the need for a multi-faceted, multi-disciplinary, holistic approach to advancing the state of knowledge on cybersafety in online social networks, and the ways in which it can be researched and protected. The development of cutting-edge research agenda and technical roadmap is needed, that will allow the community to develop and embed tools to detect malice within the systems themselves and to design effective ways to enhance their safety online.

The virtualized infrastructure of cloud computing enables many mutually distrusting users to be co-hosted on the same physical hardware. As such, it is imperative that the virtual machine manager (VMM) securely isolate users from one another. Recent research on side channel attacks have shown, despite best efforts, that the isolation and compartment guarantees offered by current VMMs are violable. These attacks are increasingly relevant because the cloud computing market shows no signs of shrinking and, in fact, will grow exponentially as Fog computing matures alongside IoT. Therefore, innovative research is needed to strengthen the security of hypervisors and VMMs against side-channel attacks by: (i) stress-testing current VMM solutions against known and new side-channels and developing common security criteria for virtualization framework assessment; (ii) offering a hardened VMM implementation on Intel based cloud computing system providing process isolation and compartment against side-channel attacks.

7. Machine Learning, Artificial Intelligence and Autonomy

Autonomous systems are a key factor in offsetting the scale of assets by an adversary, as well as performing missions that are outside the scope of traditional manned assets. (Semi-)Autonomous systems are likely to operate in contested environments in potential future conflicts. Such contested environments will contain complex dynamics and uncertain operating conditions, for which adaptive/learning methods can be applied. This section summarizes adaptive/learning methods topics that were discussed during a series of breakout sessions focused on Machine Learning, Artificial Intelligence and Autonomy.

A research agenda for artificial intelligence and autonomy

Machine Learning, in particular Deep Learning (DL), has been heralded as the solution to create intelligent, autonomous machines. Unquestionably the recent advances of DL are exceptional, but they require further improvement before they can be used to create reliable engineering systems. DL capitalizes on the theory developed in the second golden age of neural networks in the late 90's, the existence of huge databases and cloud computing to solve practical problems that are of interest to the industry, such as speech recognition, machine-based text understanding and a restrictive content based image retrieval. The current belief by some communities is that the framework and models of DL can be applied to any problem, which may be correct or simply the “first step fallacy”, so common in evolving science and technology.

The popular success stories of machine learning announced at Google, Facebook and Microsoft concern learning from data in a purely static setting. An example is Alpha Zero – Google's computer chess master that was trained by allowing the computer to “play itself”, while a reinforcement learning procedure was used to learn optimal decision rules. The computational effort for training is almost beyond comprehension:

- 70 million chess-move evaluations each second.
- One day of computation on the Google server farm is equivalent to over 12,000 years of computation on the best Apple laptop.
- Total training time was about 40 days, or 138,240,000 seconds – this corresponds to close to 10^{16} chess-move evaluations (Abbeel, Schulman, Singh, & Silver, 2017).

Similarly, the success stories on facial and speech recognition and translation relies on the the availability of huge training datasets. Google's language translation engine, for instance, was made possible by training on massive datasets of human translated texts (Lanier, 2014).

In short, the current success stories of ML and AI, impressive as they are, are mainly due to the application of massive computational power and equally massive training datasets. In fact, it can be argued that apart from the availability of huge databases that models can be trained on, and the advent of enormous computational power to enable such training, there has not been much progress in the foundational science of machine learning or artificial intelligence.

To transform ML/AI into tools that solve challenging problems such as creating reliable autonomous systems, further advances will be needed in the foundational science and mathematics of these tools. This is especially true in applications of ML to control of dynamical systems.

Advancing the science of machine learning

Current DL successes share a very important characteristic: one can create a “bounding box” around the volume of possibilities, which enables the use of data as “cheap priors” that simplify the architecture of the learning machine. Indeed, recall that speech recognition has a small discrete complexity (44 different phonemes – sounds - in the English language), therefore more data will not change the number of classes that the learning machine will need to recognize. In cases like this, the more data we obtain from different persons the more points we have to fill the volume of the space, therefore the “unreasonable effectiveness of data”¹ applies. However, a formal analysis of what types of problems benefit from this solution remains undone. One hypothesis is that when more data affects the volume of the possible solutions, i.e. increases the bounding box, the unreasonable effectiveness of data disappears.

Another important class of problems is when we do not have large datasets. Many DOD applications share this characteristic (automatic target recognition, scene understanding, autonomous navigation). These class of problems are characterized by continuous changing scenarios (space-time), where some patches are similar from previous occurrences but most of the context is different. In this solution space, current learning techniques are lacking a systematic way to set up the problem. The closest class of machine learning methodologies is online learning (also called streaming data), which is currently being solved using the same basic paradigms as static machine learning, i.e. build a model that applies to the full space of the data. This can only be done statistically by incorporating hypothesis that are unrealistic (such as stationarity), or by conditioning the current data on the past. Because statistical conditioning is computationally complex, we normally just condition on the present, yielding the Markov model, which is a definite improvement but not sufficient. It is therefore not clear that a statistical approach is the right one to tackle the problem.

Instead of this approach we can seek inspiration from the cognitive brain². Biology “invented” the brain (initially a collection of ganglion neurons) in the transition from static organisms (like plants) to moving organisms. It is interesting that there is a species of sea squirts that has a brain as a mobile larvae and digests it in adulthood when it attaches itself to the sea floor. We can say that brains exist to improve survival of a moving organism in an unpredictable world. Therefore we have to understand the role of each brain component in a 4-D space time world. Brains cannot afford to find expensive optimal solutions: they provide the best solution within a prescribed amount of time and limited by the available computational resources. We know that a lot of the processing is memory based, i.e. biological organisms extract important events from the continuous flow of time, store them and then utilize effectively the information for the current situation. The brain is a heterogeneous bidirectional (bottom up and top down) architecture of different processing modules, each dedicated to a specific task, coordinated by expectation and reward mechanism down in the basal ganglia, and driven by the sensory input.

An example of such a cognitive architecture for object recognition in video, which attempts to explain the input using redundancy reduction approach (unsupervised learning), is in (Principe & Chalasani, 2014) . This architecture is called the deep predictive coding network

¹ Halevy A., Norvig P., and Pereira F., “The unreasonable effectiveness of data”, IEEE Int. System Magazine, pp-8-12, March/April 2009

² Fuster J., “Memory in the Cerebral Cortex”, MIT Press 1995

(DPCN). The DPCN is a type of DL learning architecture, but only implements the “visual cortex”, which shows that the current paradigm of DL is missing the point that perception is not a goal in itself, but simply a module of the overall architecture for artificial intelligence. We are now coupling the DPCN with memory modules, to extract automatically object properties that are not contained in the “pixels”, i.e. they are internal constructs that the system can learn in confabulation mode (i.e. even without new inputs the system can find new relationships between its stored internal representations). This system does not forget and can continuous learn new cause-effect relations and assign value through reinforcement learning.

From the architecture discussion, it is obvious that the current algorithmic tool-set is insufficient. For instance, the way one handles memory in neural networks is by using recurrency, i.e. memory in the state, which has a fundamental compromise between depth and resolution. We do not have efficient implementations of how to combine external memories with ongoing computations conducive to self-learning. Our networks also forget the past when they learn new events. Most of our learning techniques are still limited to quantify second order statistics in the data (like mean square error), we lack tools to understand the learning process which should be based on information theory, and there is still no practical approaches of quantifying generalization, which is a major drawback. In fact, machine learning methods are impractical because it is not possible to calculate an envelope of performance around the current operating point, as we do in engineering systems. Until this is achieved they cannot be deployed with guarantees of performance.

Another issue is the role of memory access when dealing with large datasets for machine learning since the data may not even fit in a hard drive let alone in processor memory. Algorithms that are tailored to work effectively with distributed and remotely located data will benefit machine learning advances (Pardalos, Conca, Giuffrida, & Nicosia, 2015). This is particularly true for real-time control of dynamic systems using learning techniques for control computation. In such applications, even though the data sets in question may not be as big as in static learning applications, they are likely to be large enough to still suffer from some of the same issues. In addition, speed is of special concern in control, making the problem more challenging and imperative.

A common class of data that current machine learning methods do not handle very well are manifold-valued data. These arise in many applications of machine learning, such as computer vision for navigation, medical imaging, etc. They include linear sub-spaces, covariance descriptors, structure tensors, shape densities etc. Computing statistics of these manifold-valued data is fundamental to many machine learning applications. Further, most of these data sets naturally occur in very high dimensional spaces but yet lie on a lower dimensional manifold of known geometry/structure. These manifolds are not linear vector-spaces but curved spaces. Thus any analysis of such data must respect the intrinsic geometry of such spaces and utilize only the natural operations defined in such spaces. This calls for borrowing tools from Differential Geometry and Probability theory in order to develop computational methods for statistical analysis of data residing on such manifolds.

In this era of massive and continuous streaming data, where data samples are given incrementally, from a computational and storage (memory) perspective, it is natural to seek

online algorithms that can recursively update the computed intrinsic statistics rather than computing it from scratch with the arrival of every new data sample. This opens up an area of computational geometric-statistics wherein, we seek to develop online algorithms to perform statistical learning and inference from very high dimensional manifold-valued data. Efficient recursive online algorithms for computation of the most fundamental and basic statistic, namely the intrinsic average/mean called the Fréchet mean of a discrete set of manifold-valued data with a variety of curvatures, has been reported in (Ho, Cheng, Salehian, & Vemuri, 2013) (Chakraborty & Vemuri, Recursive Fréchet Mean Computation on the Grassmannian and Its Applications to Computer Vision, 2015) (Salehian, Chakraborty, Ofori, & Vemuri, 2015). In autonomous navigation of vehicles equipped with heterogeneous vision sensors, techniques specially designed for manifold-valued measurements and states can provide better accuracy than competing methods that seek to use vector space methods that are informally adopted to manifolds (Knuth & Barooah, Collaborative localization with heterogeneous inter-robot measurements by Riemannian optimization, 2013).

The mathematical challenge here lies in proving the statistical consistency of the recursive estimator. Finding the Fréchet mean, a.k.a the Riemannian barycenter is central to many Machine Learning and Computer Vision algorithms such as K-means clustering and PGA (principal geodesic analysis) which is the manifold-version of the popular PCA, subspace clustering, image segmentation, etc. Moving forward, in this context, it will be useful to tackle the fundamental issue of statistical robustness (to outliers) i.e., to develop recursive online algorithms to compute robust statistics and perform statistical inference from these high-dimensional manifold-valued data sets. Developing theory and efficient algorithms for such problems will be apt for streaming data.

Machine learning for autonomous control of dynamical systems

The remarkable success of big-data and big-computation enabled machine learning over recent years suggests that these ideas are valuable for learning and control of dynamical systems. However, specialized needs for control require specialized learning techniques. Current success stories are obtained by leveraging enormous computation power and training. A control system that is supposed to deliver autonomy cannot accommodate the delays required for such computations on parallel super computers to obtain a single control solution.

Because of the requirement for a real-time “good enough” answer rather than an optimal answer that takes far too long to compute, the use of additional structures may be necessary. Based on a model or simulator of a physical system, we require a family of candidate control solutions that can be tested and compared in real-time in realistic computational environments. This requires constraints on the complexity of the policies, and also time-limitations on computation.

A direct approach is tractable in some cases: given information constraints and constraints on the decision rule (such as linearity and complexity constraints), it is possible to apply actor-critic methods to obtain the best policy within this class. A-prior structure that is not sensible in chess is perfectly reasonable in applications to flight control.

The essential philosophical starting point of control systems engineering is not likely to change: a model or surrogate “prior” on system behavior is constructed based on physics or observations of the system to be controlled. This can then be used to create learning algorithms tailored to the particular application. For example, in optimal control there is an associated value function for which a great deal of structure may be known: convexity, rate of growth, or even a parametric form (Bertsekas & Tsitsiklis., 1996) (Borkar & Meyn., The ODE method for convergence of stochastic approximation and reinforcement learning, 2000). This is just the right kind of structure that can be used to design specialized algorithms based on TD- or Q-learning (Devraj & Meyn, Differential TD learning for value function approximation, 2016; Devraj & Meyn, Zap Q-learning, 2017; Devraj & Meyn, Fastest convergence for Q-learning, 2017). Additional methods exist that simultaneously learn a model, while also developing an approximate solution (Kamalapurkar, Andrews, Walters, & Dixon, 2017) (Kamalapurkar, Walters, Rosenfeld, & Dixon, 2018) (Deptula, Rosenfeld, Kamalapurkar, & Dixon, 2018). Motivation also exists to accelerate learning by learning off-policy (i.e., learning about actions in a separate part of the state space while executing). Concepts like Bellman Error extrapolation (similar to TD-learning) can be used to efficiently learn off-policy. (Kamalapurkar, Walters, Rosenfeld, & Dixon, 2018)

Verifiability is often referred to as one of the most challenging barriers toward the development of autonomous systems. Verifying modern “non-autonomous” software is already a daunting task that requires extensive amounts of costly and time-consuming testing. For autonomous systems, additional challenges come from the fact that they are expected to function in a variety of operational environments whose full characteristics are not known a priori. Indeed, the Air Force's Technology Horizons report assessed that the anticipated impact would depend on developing “entirely new methods for enabling trust in autonomy through verification and validation” (Dahm, 2010).

Integration of autonomous systems at scale hinges on factors including (i) how capable they are in delivering complicated missions in dynamic and a priori unknown environments and (ii) how we can establish trust that they will operate safely and correctly. These two factors are partly in conflict: The former requires adapting to and learning new skills in unpredictable, dynamic environments. However, for the latter, we lack the means to formally characterize and systematically design for provably safe and correct behavior of such adapting and learning systems. The current practice based on testing is not sustainable due to the complexity of such systems, which stems in part from the excessive amount of operational and contingency scenarios they are expected to handle.

A design artifact is called *provably correct* if an appropriate abstraction of it can be proven, by design or by post-design verification, to satisfy given formal specifications. Expressing unambiguous specifications of the operational requirements and safety constraints an autonomous system is to satisfy, the kind of environments in which it is to operate, and the sensing and communication capabilities available to it is an essential yet often overlooked first step of any design flow. A design methodology in which, given such formal specifications and mathematical abstractions of the system, design artifacts are automatically *synthesized* and compiled into executable control protocols is increasingly becoming a necessity. In such a methodology, the correctness of the resulting protocols is granted the integrity of the

specifications themselves due to the proven properties of the underlying compilation algorithm. Hence, it is often referred to as *verifiably correct by design*.

The value of *verifiability by design* will only be amplified as the emerging capabilities for *learning-based adaptation* are incorporated. In a learning setting, the intent of the designer is typically encoded in a mathematical construct, called a *reward function* (a scalar-valued function over the relevant features, states and inputs of the system). While it may be tempting to design the reward functions used in learning in order to “program” the system to operate in a desired manner, reward functions intrinsically lack the expressivity necessary to reason about complex requirements for safety-critical autonomous systems. Moreover, in their current form and functionality, they are difficult for human designers and operators to work with as a way of reliably specifying tasks. Furthermore, reward values are often dependent on or specialized for particular environments. This dependence requires tedious adjustment for each new target environment, undermining the very purpose of learning. Finally, over-engineering the reward structures to encode complex requirements runs the risk of triggering negative side effects or hiding potential bugs that surface only when the learning agents are to be integrated into broader autonomous system operation.

While the conventional abstractions for learning have served well in static, isolated applications with significant domain understanding, they fall short in dynamical and adversarial settings. Existing learning abstractions and algorithms are often oblivious to the need for verifiability. Additionally, the majority of DoD applications for autonomy are *data-starved*. We argue that the increasing success in learning without an explicit effort for verifiability will hinder its transition to real systems and limit its impact.

A case study exemplifying the need for a platform that integrates autonomy services (e.g., data-driven learning) and model/specification-based synthesis is provided in (Fu & Topcu., 2014). A *run-time shield* that enforces specifications including collision avoidance is synthesized and applied to an autonomous system. Run-time shielding with respect to safety specifications not only ensures safety during learning and execution but also improves data efficiency in learning by restricting the space for data-driven exploration and focusing on the behavior that is plausible for the mission. Additional examples of recent work that bridges formal methods, learning theory and controls toward developing adaptable yet provably correct autonomous systems include (Wen & Topcu, Probably approximately correct learning in stochastic games with temporal logic specifications, 2016; Wen, Papusha, & Topcu, Learning from demonstrations with high-level side information, 2017).

In a networked or multi-agent setting all the issues mentioned above get more severe. Here we use the phrase “network setting” to cover many types of distributed measurement and control situations. The simplest is that of a multi-agent situation: there are more than one agents that need to make decision and furthermore those decisions need to be coordinated in some sense. The relevant information and measurements are spatially distributed, and no agent has access to all the relevant information. Even for a single agent such spatially distributed complex measurement patterns may arise.

Synthesis of decision rules subject to distributed or complex non-local measurement constraints is quite challenging and no effective formalism currently exists to solve these problems. The state-of-the-art is based on distributed consensus-type algorithms that work as a mechanism for information diffusion. However, the convergence rate of such diffusion-based schemes are slow and uncertain, and have not seen much translation to practice. Novel approaches are needed that satisfy the need for speed, reliability and robustness.

Asymptotic theory such as mean-field games can be used to obtain insight on the structure of decentralized control schemes in a homogeneous distributed environment. This can again be used to obtain learning architectures to obtain the best control solution within a given family (Mehta & Meyn, 2009) (Meyn, 2007). In some cyber-physical networks, there are locally measurable signals that provide global information about the aggregate state of the networks. Examples include frequency and voltage in power grids. When such signals are available, they can be used for coordination among agents without inter-agent communication, which increases both speed of convergence and robustness to malicious cyber attacks and benign failures (Brooks & Barooah, 2016).

It should be emphasized that for decision-making, learning is required on many timescales: partially observed states that may evolve rapidly with time, model parameters, and learning of gains for both state estimation and control.

Many modern-day air force operations comprise of multiple interdependent autonomous missions that are specified and coordinated in real-time by human operators. These missions evolve concurrently and can be highly dynamic as the environment and task specifications can change rapidly with time. The presence of increasing levels of mission autonomy and the dependence on unreliable communication for mission coordination, has made mission planning increasingly vulnerable to cyberwarfare techniques (e.g., sensor spoofing, communication attacks). Recently, several high-profile attacks on autonomous systems have been reported, including the highly publicized capturing of the RQ-170 Sentinel US unmanned aircraft (using a combination of GPS spoofing and DoS communication attacks). One of the reasons for the current vulnerable state is that resiliency against attacks has not been considered as a design requirement for mission planning in conjunction with realistic, unreliable, communication models. Consequently, *addressing resiliency against various cyber and physical attacks under realistic communication scenarios is necessary to ensure successful mission completion in hostile environments, even in the presence of attacks that can critically hamper mission capabilities, while providing accurate and up-to-the-minute situational awareness to the human operators that is required for mission planning.* For this to become possible, a distributed control, estimation and communication system needs to be formed between the missions and the operators that: (i) is resilient to compromised missions and information exchanges; and (ii) can achieve coordination in the absence of reliable continuous communications. In fact, it is well accepted that wireless communications in austere non-line-of-sight environments, as those that are common in combat environments, are severely impacted by uncertainty and are highly intermittent. At the same time, resiliency in distributed control and estimation processes can handle compromised components agents (e.g., robots or sensors) and communication links as long as there are still sufficient trusted communication channels and available observations that can be used to detect and isolate any malicious agents. Therefore, a great challenge in obtaining

resilient mission planning strategies is the development of attack detection and secure estimation techniques in the presence of intermittent communication. The goal is to develop resilient mission planning strategies in the presence of unreliable and intermittent communication. In relevant literature, distributed control, optimization, and estimation methods in the presence of unreliable and intermittent communication always assume that connectivity will always eventually be regained. In practice, however, this is not a reasonable assumption. Instead, control frameworks that guarantee intermittent network connectivity are much preferred.

Advancing Human-Machine Interactions & Teamwork in AI

Until true AI arrives – assuming that is both possible and desirable – we must develop AI systems that can assist humans effectively. This vision is sometimes referred to as mixed autonomy. In many DoD applications mixed autonomy is preferable over full autonomy, even when full autonomy is possible, due to legal reasons.

To optimize human-machine teamwork, we must determine the foundational science that defines teamwork, such that we can do better at predicting team outcomes. We must also define the measures and metrics of teamwork, and isolate team performance from overall performance such that the metrics will indicate when a team worked well together regardless of the team's overall task performance. Finally, determine what (if any) are some useful and repeatable team patterns.

Humans are very good at solving motion planning problems. Using experience and a map of the world, they can quickly find a path between a starting and a goal location that is also optimal in some sense. Nevertheless, the ability of humans to plan is typically limited to low-dimensional spaces and the resulting plans are high-level, e.g., sequences of regions to visit rather than continuous trajectories. On the other hand, machines have the ability to solve complex planning problems in continuous high-dimensional spaces. However, optimality comes at the cost of computational complexity so that oftentimes only a feasible solution is possible. Moreover, machines typically treat every problem instance as new and do not use experience. Future motion planning algorithms will need to combine (i) human experience, context-awareness and ability for high-level reasoning with (ii) the ability of modern computers to plan in high-dimensional spaces. Specifically, the goal is to develop synergistic frameworks where human experience can help with optimization when computational complexity is too high for a machine alone to solve a planning problem, and the ability of machines to plan in high-dimensional spaces can help humans refine their view of the world and shape their experience. While empirical or statistical methods for human-robot collaboration exist in the literature, no principled framework exists that comes with performance guarantees.

Human agents do not act in isolation but are influenced by both their information environment and the communities to which they belong. In the context of an agent's information environment, messages received, and longstanding individual and group preferences, an agent will update its internal state and the state of the communities to which it belongs. The outcome of the state update determines the agent's next pathway. For optimal human-machine teaming, a model of each agent's preferences and states may be useful in predicting agent actions and understanding agent goals for improved outcomes during autonomous operations.

Social researchers have shown that when individuals are in teams, particularly large crowds, their behavior changes (Altermatt & Pomerantz, 2005; Friedkin & Cook, 1990; Le Bon, 1896; Stage, Dezechache, Pierre, & Grèzes, 2015; Turner & Killian, 1987); specifically, when individuals join groups, they promote their group-mindedness while demoting their individuality. Theories have emerged to explain this phenomenon (e.g., contagion theory and emergent norm theory (Le Bon, 1896; Turner & Killian, 1987; Park & Turner, 1967); the suitability of these theories for modeling and predicting behavior within the context of human-machine teams is unknown but there are a few reasons to believe that the same conceptual relationships that explain human-human teamwork in physically proximate situations can be applied to human-machine teamwork and autonomy.

Several differences between human-human and human-machine interactions are crucial. In particular, existing human-machine interactions are limited by artificial interfaces such as computer screens, keyboards, or at best speech interfaces; human-human interactions have no such limitations and can take advantage of multiple modalities of interaction information such as gaze, facial expressions, and tone. However, technology is advancing to the point that these additional modalities could be incorporated in human-machine interactions.

Many human-human teaming models assume that agents actively participate in the process at all time steps, for a pre-defined time span, and in similar ways across time. However, this is hardly true in today's online information environments, where asynchronous and irregular communication is the norm. Thus, accurate models of teaming must include an analysis of behavior over time, such as communication time or frequency, or message topic or sentiment, and detect any anomalies. Additionally, as online groups, communities, and factions have been shown (Nguyen, Boyd-Graber, Resnik, & Miler, 2015; Mitchell, Hollingshead, & Coppersmith, 2015) to have their own code words, their own ways of framing an issue, and their own individualized reactions to current events and issues, human-machine teams may be more successful if they share a common language/communication model. Modeling this shared language and communication strategies to identify outliers and anomalous communications may enable the discovery of the most effective human-machine teamwork strategies.

It is possible for the next generation of computer user interfaces to be cognitively based through multi-modal biometric data such as eye tracking (gaze and pupil size), EEG, facial expressions analysis, tone of voice and natural language processing. The machine should be able to adapt in real time to the mental state and thought process of human agents (e.g. through a user interface with an adaptive level of fidelity and level of abstraction) and the human agents' feedback provide guidance to the machine. Such bidirectional feedback resembling human-human interaction is likely to result in better performing, more robust, better accountable, fault tolerant systems that are resistant to adversary manipulation.

There are several trends that that points to this step being possible and feasible. The “unreasonable effectiveness of data” is our ally since there is a finite set of relevant human mental states while the amount of biometric sensor data is increasing rapidly. Advances in computing ability, including cloud computing, enable multi-modal biometric sensor data analysis and integration. In many cases the cognitive load of current user interfaces e.g. for pilots, command and control and multi-platform drone swarm operators is already at its limits.

The cognitive user interface task can be tackled with a combination of formal methods that define the interactions and an applied empirically validated approach where first a prototype of a cognitive user interface is tested and after successful validation the method applied and adapted to a military relevant task. One possible platform for such prototypes is the ancient game of chess, that is so called freestyle or cyborg chess where teams consisting of both humans and machines compete. Chess could thus play a similar role in cognitive user interface development as it has played in machine learning (most recently AlphaZero) and cognitive psychology.

8. Future Research Directions in Multi-Agent Sensing and Communication

The ability of the USAF to project force can be multiplied through multi-agent systems. Fundamentally, multi-agent systems function by having a team of agents share information in order to jointly coordinate their activities, and there is significant interest in multi-agent systems within the autonomous systems community. Autonomous systems research involving multi-agent systems typically focuses on designing and analyzing coordination algorithms, with often implicit assumptions concerning their communications and operations over time. In the future, multi-agent systems are expected to operate in complex, dynamic, and contested environments, which can cause assumptions about their ability to communicate and assumptions about prevailing conditions to fail.

For these and other reasons, new research is required to mitigate practical challenges facing multi-agent systems. This section summarizes multi-agent network topics that were discussed during a series of breakout sessions focused on Future Research Directions in Multi-Agent Sensing and Communication. The following describes research opportunities along four thrusts. First, we outline research opportunities in intelligently composing multi-agent systems. Second, we discuss the interaction between control and communications in such systems and the need for further developments to optimize performance across these domains. Third, we describe future developments which will allow multi-agent control systems to operate successfully under realistic and challenging conditions expected in practice. Finally, we discuss how communications technology must evolve to scale to support large multi-agent systems and present some possible directions for research.

Reformulating Network Composition

This section describes several research opportunities pertaining to selecting assets to participate in a multi-agent system. Typical approaches to multi-agent control often assume that a network is given and then focus on how to control it. However, there are opportunities to tune the network itself to be suitable for a given objective, and there is the potential to enhance network performance by selecting appropriate network members before runtime.

Research in multi-agent coordination and control is almost exclusively focused on decentralized approaches in which communications among agents are point-to-point and in which each agent communicates with some others in a network, though not all others. For example, each drone in a swarm may exchange information with other nearby drones, but not with other drones that are far away. Accordingly, algorithms designed for coordinating such systems must only use computations that can be executed locally by each agent using only the information it has access to. This condition can be quite restrictive, especially for operations that cannot be decentralized in general, such as matrix inversion.

Instead of insisting on pure decentralization, one could envision incorporating some centralized information into a network. To do so, a mobile agent could circulate among others to gather information from each agent and ferry it to the others, or a stationary base station could be used to gather information, process it, and broadcast the results. Regardless of the exact form, adding some centralization into a network eliminates the need to decentralize all computations. Moreover, adding some centralized information into a network provides agents with a stream of

information that would otherwise be absent. This centralized information may be shared with agents at a slower rate than other information in the network simply because it is centralized, though making additional information available to the agents will typically be useful to them. Determining if, when, and how centralization should be incorporated into a network remains an open question. The existing literature already has some purpose-built approaches to incorporating centralization into multi-agent networks, though there does not yet exist a general theory of mixing centralization and decentralization in a principled and automatic way. While existing initial developments are promising, further theoretical developments are required to develop a general-purpose theory of mixed centralized/decentralized multi-agent systems.

Perhaps the foremost driver of multi-agent control developments is *scalability*, which broadly means that a multi-agent coordination strategy should work equally well regardless of the size of a network. The ability of an algorithm to scale is regarded as a key bottleneck in the ability to use large teams of agents, though communications present an additional bottleneck as well. Fundamental limits of bandwidth present limits on the number of agents that can be communicating simultaneously, and, as a result, communications limit the number of agents that can reliably interact when working together. Communication limits depend sensitively upon the conditions encountered by teams of agents, the means of communicating, and the presence or absence of adversaries jamming communications. If these factors can be anticipated before runtime, they can be used to identify limits on the volume and quality of information that can be shared. These limits can then be used to inform the sizes of teams that are deployed.

There does not yet exist a comprehensive theory for team sizing based upon communication limits, and future research developments are required to fully account for how communicating can restrict the scale of multi-agent systems. Such a theory would presumably need to account for how communications degrade as teams scale up due to interference and other factors, and additional developments are required in both communications and control to assess how communicating affects networks' ability to scale up.

Closer Coupling of Control and Communication Strategies

Control approaches to multi-agent systems typically (and implicitly) assume that communications are perfect, in the sense that information is assumed to arrive at its destination immediately and in the exact form in which it was sent. This assumption rarely holds in practice, and communications can be asynchronous, due to latency, and noisy, due to data becoming corrupted in transit. Those control schemes that do consider latency or communication losses assume that these problems are random and unaffected by the actions of the control algorithms. However, in reality, latency and packet losses will be highly sensitive to the actions of the control systems. For instance, the locations, velocities, and orientations of autonomous vehicles will significantly change the channel conditions, network maintenance traffic, and latency in routing traffic through a network. Control applications may also alter network performance by the amount of traffic used in control and its quality-of-service requirements. Thus, new approaches need to be developed for control and communications/networking that allow both of these technologies to better adapt based on their joint set of requirements and constraints.

Agents in a network may share information of different types to facilitate multiple ongoing tasks. For example, a team of drones may exchange position information to remain in a fixed formation

as they photograph an area en route to a common destination. As a result, not all information is necessarily of equal importance. Communication limitations mean that constantly sharing all information is not feasible, and, instead, higher value information must be prioritized in communications. A single agent may gather large amounts of data from several data streams, such as video cameras, thermal cameras, and radar. Determining which data should be shared is critical to avoid congesting available bandwidth, and new theoretical developments are required to determine how to prioritize communications in this way. A key challenge in doing so is predicting which information will be most useful to other agents. Although agents may be collaborating on some task, no agent will generally know all information available to other agents. As a result, an agent cannot be certain of which information is most needed by another agent, and some form of heuristic may be required to predict this need. Jointly with prioritizing information, agents will also need to determine the fidelity with which information must be sent. There exists a broad range of possible ways to share information, and an agent must be able to predict how much noise is tolerable, how accurate data must be, and other factors that affect how data is encoded and shared. Accordingly, additional developments are required to understand how the end use of a piece of data should affect how it is shared.

Asynchronous communications arise in networks for several reasons, including when facing anti-access/area denial (A2AD) measures, when networks are large, and agents lack a common clock, and when agents communicate over large distances. In all these cases, asynchrony means that agents will receive different pieces of information at different points in time. Agents may therefore have disagreeing values about all pieces of data, such as the location of a target or the state of another agent in the network, though they must work together despite these disagreements. Existing multi-agent coordination strategies often assume that communication delays are bounded, or that certain graph-theoretic connectivity properties are satisfied over time. However, agents will often be unable to enforce these assumptions because their communications are affected by a wide variety of factors beyond their control. In addition, agents will often lack a common clock and will therefore be unable to even verify the satisfaction of these assumptions. For these reasons, further algorithmic developments are required to mitigate the effects of asynchrony in networks without imposing assumptions to limit the degree of asynchrony.

In wireless communication systems, tradeoffs exist between reliability/fidelity and latency. For instance, data that experiences errors during transmission can be used immediately after reception, but the data is corrupted. However, automatic repeat request (ARQ) can be used to request retransmission of corrupted data, at the cost of additional latency. Similarly, error-correcting codes can be used to protect data from corruption, but a high level of protection requires low code rates, which increases latency and reduces throughput, whereas higher code rates have lower latency, higher throughput, and a greater chance of data corruption. Sophisticated control and communication designs can tradeoff among these factors to maximize the performance of a multi-agent system. For instance, agents may decide whether to use an error-tolerant control system information with noisy information or use a delay-tolerant control system and wait for noise-free information. Most interesting is that the amount of corruption and delay can be estimated in real time and new control approaches may be developed to adaptively switch among different operation modes.

The requirements for accuracy may also change throughout a mission. For example, a drone swarm tracking a target may be able to tolerate noisy data sharing when far from their target, though they may wish to share exact coordinates as they approach it. In this case, noisy data sharing may be tolerable initially, though noise-free data is required later. New research developments are needed to understand how to trade off fast and noisy information with delayed and noise-free information. Doing so will reduce bandwidth usage and will accelerate some multi-agent coordination algorithms by not requiring them to wait for corrupted data to be re-sent.

If noisy data is used as input to a control system, the effects of these data will be propagated forward in time to future timesteps, and thus a single instance of noise or latency will have a lasting influence. These effects can build up over time and can substantially degrade performance in some cases. To mitigate their effects, controllers must be designed that tolerate imperfect information entering feedback loops and that can provide information to the network on when high-fidelity data is required to maintain stability of the systems. Future developments are needed to understand the impact upon agents' local controllers, such as in a single drone's position control, as well as on networks of agents in aggregate, such as when they are jointly completing some task.

Developing Practical Performance Requirements

After selecting a network and accounting for its communications, the next natural question is what the network should actually be doing. Existing multi-agent control literature has developed techniques that succeed under ideal conditions, though the definition of "success" in multi-agent control and the conditions required to attain it need not always be the same. Below, we delineate three future directions for research that will broaden the capabilities of multi-agent systems, especially in uncertain, complex, and dynamic environments.

Existing multi-agent controllers are frequently designed to give full autonomy to the agents involved. However, teams of agents may be required to assist a human warfighter in some circumstances, e.g., when having a team of drones fly alongside a human pilot. Such instances require agents to be autonomous in some decisions, such as generating control inputs to maintain a fixed distance from a human pilot or another autonomous agent. However, other decisions, such as to engage a target, either cannot or should not be done autonomously and should require human input before a decision is made or an action is taken. The needs of autonomous agents may also change over the course of a mission; one could envision having a drone autonomously fly to a particular destination, autonomously observe a target, and then recommend engaging or not engaging a target. New research developments are needed to develop controllers that permit a spectrum of autonomy that can be changed online and in real-time so that autonomous agents are always best assisting human warfighters.

And the notion of recommending an action reveals further benefits of autonomy that require additional research. Recommending an action to a human warfighter requires the warfighter to make a decision, and having too many agents recommending actions simultaneously can unnecessarily encumber the warfighter. Simply determining when an autonomous agent should recommend an action is a key research question that must be addressed, and this work should incorporate not only limitations on which actions an

autonomous agent should enact on its own, but also the desire to limit the number of decisions a warfighter must make.

Current approaches to multi-agent control typically state assumptions upfront that are required to hold and then predicate the success of their controllers upon the constant satisfaction of these assumptions. Common assumptions range from bounded information delays and graph-theoretic connectivity properties to convexity and immobility of obstacles in an autonomous agent's workspace. However, it is easy to envision these assumptions being violated, especially in contested, complex, and dynamic environments that autonomous multi-agent systems are expected to face. Violating these assumptions typically results in the outright failure of the multi-agent control process, and their frequent violation necessarily implies frequent failure. To prevent the failure of multi-agent systems, further research developments are required to promote graceful degradation of performance of multi-agent controllers when prevailing conditions do not meet standing assumptions.

One method of doing so is to have controllers that "almost" succeed when assumptions are "almost" met. Allowing degrees of satisfaction for assumptions naturally leads to the notion of a parametric assumption, i.e., an assumption which relies on some mathematical parameter falling in some range. For example, the velocity of a target may be assumed to be bounded within a specific range, or the length of a communication delay may be assumed to fall in a certain range. Degrees of constraint violation can then be studied by examining the extent to which parameters fail to satisfy the specified bounds. These degrees of violation can then be used to define degrees of success of a controller, and, in particular, the magnitude of violation can be made to correspond to the magnitude of failure of a controller. Further theoretical developments are required to realize this idea and to use it to ensure that controllers' performance degrades gracefully rather than having outright failure based on binary assumptions.

In some cases, partial failure of a multi-agent system cannot be avoided. For example, if one drone in a team suffers an irreparable mechanical failure, then certainly there is no algorithmic development that can reverse this type of failure. On the other hand, the effects of partial network failure could be mitigated by some appropriate coordination strategy among the agents. For example, formations could be rearranged to account for the absence of a member of the network and surveillance operations could be adapted to require each agent to increase the area it covers. Other failures are also possible, such as total loss of communications among agents due to mechanical failure or jamming. Existing approaches often abort any mission in progress in the presence of any partial failure of this kind, though it may still be possible to perform useful work and attain partial objective satisfaction even in the face of such failures. Additional theoretical developments are required to develop useful failure modes for multi-agent systems, especially in autonomous operations.

Evolving Communications to Support Large Multi-Agent Systems

Current wireless communication systems have limited ability to scale to large distributed systems, such as those we envision will be desired for future multi-agent missions. Existing systems suffer from the conventional approach of pre-allocated, stovepipe communication and one-size fits all networking schemes. New approaches must be highly adaptive and cognitive in the use of spectrum, able to optimize communications parameters based on mission priorities,

and able to operate effectively across heterogeneous platforms. Although cognitive radios and dynamic spectrum access has been studied for many years, many fundamental issues remain unresolved, and these are exacerbated when applied to military communications.

An important issue in all distributed systems is determining what information should be exchanged and how often. In wireless communications, we must further ask what latencies and loss rates can be tolerated. Back-channels may facilitate distributed spectrum use among users, but how should conflicts be arbitrated? How should we handle situations when the back-channel is disrupted? How should such a network operate when completely disconnected from other networks, for instance, due to jamming or interference? An important consideration is determining the most important features to adapt and adaptation algorithms that are simple and robust to new operating conditions.

Other approaches to scaling communication and coordination among large teams of agents may be necessary. In some cases, traditional RF links may be replaced using other sensors. For instance, vision can be used for location updates in place of wireless location updates in some cases. Another possible approach for unmanned vehicles is to use patterns of motion of the vehicle to convey information to neighbors that detect that motion via vision. Distributed wireless networks often become dominated by network coordination traffic when the scale to even tens of nodes. New approaches to allow networks to function with only minimal coordination information are needed. For instance, agents may use Kalman filters to predict where other agents will be, and network updates are only sent when an agent knows that his peers' predictions are bad. More generally, new science is needed on how to choose network topologies for multi-agent systems, including assessing the costs and benefits of different distributed network topologies.

Electronic warfare (EW) is often treated separately from communications, but multi-agent systems with wireless communications may offer new approaches to EW. For instance, many small, low-cost mobile agents can be positioned close to a receiver or network to jam communications or perform eavesdropping. In some cases, multiple agents may be able to be positioned in such a way that beamforming can be used for jamming, eavesdropping, or even forming high-speed backhaul links.

9. Space Systems

Space is becoming a more crowded and adversarial environment, dictating the need for increased resilience. To provide such resilience in the unique environment of space, new breakthroughs are required along several fronts. This section summarizes topics that were discussed during a series of breakout sessions focused on Space Systems. In the following sections proposed advancements in three key technology areas are provided: Space Situational Awareness (SSA), Small Satellite Technology (SST) and Autonomous Space Systems Guidance, Navigation and Control and Optimization (AO). The following sections discuss the challenges in each area and the arising basic research opportunities, along with recommendations for the Air Force.

Space Situational Awareness

Space situational awareness (SSA) and what has come to be known as space traffic management (STM) are part of a larger discussion on sustainability, currently underway in the United States and in international fora such as the United Nations Committee on the Peaceful Uses of Outer Space. In the US, sustainability is featured throughout the 2010 National Space Policy and the 2017 Space Policy Directive 1, establishing a US commitment to revisit the Moon by 2025 as a precursor to activities around and on Mars further in the future. The balance of the policy remains unchanged.

In these space policy documents, sustainability is discussed in terms of space activities and applications that make terrestrial life more sustainable, by way of disaster management and search and rescue capabilities, and management of natural resources. The policy goes further, however, acknowledging our national interest in responsible space conduct and the sustainability of space. This includes sustainable access to orbit. Most significant is the policy commitment to develop space situational awareness competence with participation from commercial, civil, and national security sources in furtherance of long-term sustainability of the space environment.

SSA Challenges

- Orbits are not monitored in a persistent manner. Different agencies and institutions use different sensors.
- A better and unified understanding of the space environment is needed for future orbital congestion and collision forecasting. The environment can also produce hazardous conditions that affect RSOs.
- There is a high computational cost (combinatorial) for multiple objects tracking and association.
- Sensor Tasking is also very challenging due to limitation on available resources.
- Collision and debris can rapidly take away space viability.

SSA Opportunities

- A unified framework for modeling and simulation of the space environment should be developed. Reconciliation shall include model assumptions resulting in consistent analysis, e.g. probability of collision. Multi University/Center of Excellence efforts are encouraged.

In particular, the Air Force could fund an Institute for Astronautics Research based upon Cooperative Agreements as the funding mechanism, where these could be re-bid every 5 years and assume that 8-12 research institutions could be members at any given time. This could open Affiliate Membership to industry and commercial entities to maximize technology transition. We propose a model of a joint/shared IP and Affiliate Members would have encouragement on integrating this shared IP into their own business model. This Institute for Astronautics Research would need an Administrator and to be located at a government facility. For example, someone at the Air Force Research Laboratory, a Chief Scientist with no program element role or responsibility and not in direct control of funds. Someone who is in the role of scientific advisory and not program management.

- Taxonomy to categorize objects. A scientific categorization based on their characteristics and behavior in orbit is needed. The SSA community needs data, including benchmark data. We need a lake where we can all access these openly, a virtual and collaborative environment which by definition already means it should have a .edu path and not a .mil
- More mathematical and computational rigor in uncertainty representation is needed.
- Algorithms that produce reliable data association and sensor tasking solutions should be investigated. The possibility of leveraging on parallel computing capabilities is suggested.
- Leveraging on machine/deep learning for the optimization of these algorithms is suggested.

Small Satellite Technology

Until ten years ago, space was firmly in the hands of a small number of manufacturers, owners, and operators. A small number of established spacecraft manufacturers built large, sophisticated, long-lasting and expensive spacecraft for governments, telecommunications corporations, and a few remote sensing companies. Around 60 - 80 spacecraft were launched per year globally, one or two at a time. The scale of these operations made the task of monitoring launches, identifying payloads, and tracking deployed satellites manageable.

The proliferation of small satellites, in particular CubeSats, as well as inexpensive but capable spacecraft components, readily available from internet vendors, introduced new opportunities for smallsat development. Satellite launches are offered by a multitude of international ride-share brokers. Virtually everybody can now be an integrator, owner or operator of a spacecraft. Schools build and fly CubeSats as science projects. Universities use CubeSats for low-cost space research, technology demonstration, and education. Startup companies envision global remote sensing and communications networks based on large swarms of CubeSats. Government agencies use CubeSats for science, technology development, and tactical military purposes. By number, small satellites below 500 kg and CubeSats are now the primary type of spacecraft launched: 108 CubeSats were launched in 2015 and 46 in 2016.

An emerging global trend is to launch small satellites in swarms or constellations. Global corporations are currently planning the deployment of so-called “mega-constellations” in Low Earth Orbit (LEO): Samsung plans a network of 4,600 satellites, SpaceX 4,425, Boeing 2,956, OneWeb 900, with a number of other companies envisioning networks that are smaller but still number multiple hundred spacecraft. In addition, small satellites are commonly deployed en masse, mostly on Russian and Indian launch vehicles such as the Dnepr or the PLSV. In 2017, a PSLV successfully deployed a record 104 satellites on a single launch.

Most small satellites lack collision avoidance maneuvering capability and pose a hard-to-detect collision risk for all other satellites with similar orbital altitudes, including expensive, irreplaceable national security space assets. Every CubeSat deployed increases the risk of impact damage to these assets. The deployment of mega-constellations exacerbates this problem over the next decade. Therefore, new efforts are required to track space debris objects, propagate their orbits, perform conjunction analyses, and then execute fuel-expensive collision avoidance maneuvers.

A low-observable small satellite equipped with relative navigation sensors, a propulsion system, and unsophisticated tools is a major threat to all U.S. space assets. Such satellites would be hard to distinguish from space debris, could clandestinely sneak up on space assets, and disable those using drills, wire cutters, or simply spray paint. Such attacks would not generate any additional debris except for the disabled spacecraft and would also be deniable by the aggressor.

A major question that needs to be answered to ensure space dominance is how to protect the sophisticated, big, expensive and hard-to-replace strategic space assets critical to U.S. national security. The proposed approach is to explore other domains to find a perfect analog: the

carrier battle group. The Navy uses escort ships and aircraft to create a defense bubble around aircraft carriers. Nothing enters this bubble without being positively identified. A similar approach could be adopted to protect strategic space assets. Several small escort satellites could be launched “piggy-back” with any major spacecraft. These escort satellites will be equipped with sensor suites, a propulsion system, and robotic tools. They will thus be capable of patrolling the local space around their “flagship”, identify any object entering the defense zone, contain debris objects, and neutralize potentially hostile objects. Furthermore, these escorts can also be used to perform health monitoring of the flagship and execute small repairs to mitigate damages from debris or meteoroid impacts. The flagship would be equipped with docking ports to refuel and recharge the escorts as needed.




Figure 7: Cubesats Escort Concept

Escort satellites would thus make major space assets defensible. But large spacecraft are still limited in their maneuverability and in the predictability of their orbital trajectories. To overcome those limitations, space tugs would rendezvous with space assets on demand and maneuver them to other orbit altitudes or orbital planes. The space tugs themselves would then be refueled from in-space depots. This capability would significantly increase the operational flexibility and adaptability of the national security space infrastructure.

The third element of the concept is space-based replenishment of the escort satellites. The escort satellites are not expected to have the same 15+ year lifetime as major assets. Therefore, there will be a need to replace them frequently. Orbital FabLabs could use additive manufacturing to produce simple expendable small satellites on demand for low cost and without highly visible launches. The maximum utilization of this in-situ production idea would be to actually “mine” space debris as natural resource to turn the space-grade materials in expended upper stages and discarded satellites into raw materials for the fabrication of new space assets. This would give the U.S. the capability of clandestinely adding to the existing space infrastructure.

In optical communication, ranging and constellations, several research thrusts can be explored. Optical links offer the potential for high accuracy orbit determination, navigation and high bandwidth communications for satellites that are size and power constrained. One of the key performance drivers for these systems is timing. Nearly all space navigation systems rely on

light-travel time calculations, and data throughput for free-space optical communications systems depend on the timing resolution with which laser pulses can be emitted and detected. Two research activities develop and flight-demonstrate compact precision optical technologies for space navigation and communications are proposed.

The first project is CubeSat demonstration of ground-to-space clock synchronization at the 200 ps level (6 cm light travel time) using short optical pulses. This technology could enable future small satellite navigation systems beyond the GPS sphere, for example around the Moon or Mars. Borrowing some technologies from previous optical time transfer missions, we have developed a new compact, low-power Optical Precision Time-transfer Instrument (OPTI) that will simplify the process of correcting for clock drift on spacecraft. OPTI consists of a Chip Scale Atomic Clock (CSAC), an avalanche photo detection system, a picosecond event timer, and a retroreflector.

The second project will focus on the development of Miniature Optical Communication Transmitter (MOCT), a compact, low power, optical communication system for power-constrained small satellites. The key technology behind this instrument is a software-defined pulse modulator that implements a programmable pulse position modulation scheme by generating precision delays between pulses rather than using a high-speed slot clock. The modulator is entirely implemented in a Field Programmable Gate Array (FPGA) to keep power, volume and complexity under control. A CSAC provides a reference frequency to the FPGA to generate the delays. Using a delay-locked loop to compensate for environmental effects on the modulator, a total timing error of < 10 ps has been demonstrated inside the FPGA. The modulator provides rising edge triggers for a 1550 nm seed laser diode driver, which feeds a fiber amplifier to produce the emitted high-energy laser pulses.





Figure 8: Miniature Optical Comms Transceiver

SWaP, high-speed, radiation-tolerant processors and data interfaces, secure high-speed communications systems between terminals that have only limited volume for high-gain apertures, and navigation and timing systems independent of GPS but that can achieve better than one-tenth of a wavelength ranging. Such fleets of small spacecraft will require autonomous planning and scheduling systems that can manage rigorous resource management and management of intersatellite links, and that are capable of quick replanning for hundreds of spacecraft. New spacecraft designs that miniaturize functional elements (e.g., system-on-a-chip)

can enable production to become more reliable and assembly-line like, making it possible for hundreds of nodes to be deployed and coordinated.



G. Lordos, MIT

Figure 9: Swarm of swarms

The aforementioned strategies push the envelope in basic scientific research in space systems on many fronts. For this technology to be accessible, research should be conducted in real-time GNC algorithms, real-time uncertainty quantification and robust control implementation, sensor development to detect and identify approaching spacecraft, could be vision-based, constellation design revolving around the protection of a high value space asset, space tug technologies, in orbit manufacturing, recycling of spacecraft, new efficient communication algorithms and hardware and new spacecraft designs that take advantage of constellations and redundancies.

SST Challenges

- Dependency on space services and no ability to protect space assets.
- Lack of repository for design-build-fly activities (the dos and the don'ts).
- SmallSats are hard to identify and track.
- Thousands of SmallSats will be deployed soon, e.g. OneWeb, SpaceX, Boeing, etc.

SST Opportunities

- Escort SmallSats can be used for threat identification, defense and even for service and repair for space assets. This would translate in autonomous GNC, multi-agent planning/coordination, low SWaP sensor systems and processors.

- Optimal power management algorithms.
- SmallSats can be used in swarms for surveillance, communications, alternative to GPS. Research in Laser communication, Laser ranging.
- Manufacturing of SmallSats on-orbit and even from recycled decommissioned spacecraft. Research in on-orbit manufacturing and resources harvesting.
- Uncertainty-based systems engineering by producing low-cost designs while increasing risk tolerance.

10. Unmanned Air Vehicles

The topic of unmanned air vehicles (UAVs) otherwise known as unmanned aircraft systems (UAS) or remotely piloted air vehicles is an application or the integration of most of the technologies addressed in the summit. The breadth of these systems cover numerous scales ranging from large systems such as the current Global Hawk to small systems such the Black Hornet. They also are utilized in an extensive repertoire of missions, some of which include; surveillance, search and rescue, combat, air support, strike, and communication. They can be used as a single asset or work in cooperation with multiple and/or diverse systems to add redundancy to a mission or enhance coverage. The applications are ever increasing and represent one of the fastest growing technologies in the Air Force today. Although the current systems that are utilized by the warfighter are advanced, improvements and breakthroughs in technology are paramount for the United States to remain one step ahead of our adversaries. Future systems will need to be faster, more maneuverable, stealthy, smarter, and more versatile.

This section summarizes topics that were discussed during a series of breakout sessions focused on UAS. In the following sections proposed advancements in five key technology areas of UAS will be addressed; *aerodynamics, structures, propulsion, controls, and mission support*. In each of these key areas future research and innovation are suggested. There may be significant overlap in the suggested topics amongst the five focus areas. The panel convened during the Unmanned Systems breakout sessions mostly have expertise in small UAV systems and since the technologies that drive the advancements in named vehicles closely mirrors that of large UAVs, the suggested research areas mostly focus on small UAVs.

Aerodynamics

Developing UAS that can take-off from land from constrained environments is an important feature for next generation aircraft. This ability will allow future UAS to operate in austere environments with little mission support. This can offer a significant advantage over the adversary in a combat environment.

The ability of UAS to take-off and land from constrained environments can be roughly split into short take-off and landing (STOL) and vertical take-off and landing (VTOL). The problem with V/STOL aircraft is that they often pay a penalty in terms of their performance outside the take-off and landing phase of flight. However, advances in multi-disciplinary optimization techniques (MDO) and novel propulsive and flow control schemes like distributed electric propulsion (DEP) could help alleviate these problems. Thus, future V/STOL aircraft may not have as many compromises as their current generation counterparts in terms of overall flight performance. The overall trend of aircraft development has always been in the increase of operational envelope and capability. Emerging technology could now open the pathway to mainstream adoption of V/STOL aircraft and emphasis should be given to these designs.

A significant factor in the development of stealthy UAS is the minimization of their acoustic signature. Current research areas are in the use of acoustic detection and tracking of UAS for applications in counter UAS and sense-and-avoid. Additionally, research has shown that some UAS systems can be detected over half a kilometer away. Passive detection at these ranges could conceivably compromise the mission, depending on specific mission objectives.

Thus, it seems evident that future UAS systems should be designed with minimal acoustic signatures. However, with current technology, minimization of these acoustic signatures naturally leads to compromises in performance (such as endurance, payload capacity, etc.). Due to trade-offs between acoustic signature and overall aircraft flight performance, clear acoustic / performance requirements must be set forth from the Air Force and other Department of Defense branches for future UAS aircraft.

The ability of an aircraft wing to change its shape offers numerous advantages to wings of a static design. For example, wings are often shaped for take-off and landing, and this in-turn, may make the wing inefficient for cruise flight. The ability of the wing to change its shape can potentially allow for the aircraft to be optimized for every part of the flight regime. The ability to design an aircraft with this capability is made possible through the use of morphing aircraft technology.

Morphing aircraft technology is not new, however. Much of the morphing on-board traditional aircraft occurs through the use of simple devices, such as flaps, slats and in some cases, variable wing sweep. These aircraft broadly fall into the category of “discrete” morphing aircraft. Unfortunately, to achieve anything that is more complex than simple discrete morphing adds significant penalties that have often offset any gains that comes from morphing the wing. Recent advances in smart material technologies have allowed the development of “continuous” morphing aircraft. This technology is embedded within the aircraft structure itself. Continuous morphing has distinct advantages over discrete morphing aircraft. These advantages are summarized in the table below:

Discrete Morphing	Continuous Morphing
Singular functionality	Multiple functionalities
Used locally on board the aircraft	Distributed over the body of the aircraft
Operated sparsely on the flight envelope	Operated continuously within the flight envelope
Suppresses coupling between the aircraft axes	Exploits couplings in morphing schedules and between the aircraft axes

The ability to use the morphing technology across multiple functionalities can be expressed as using the same actuator to affect the flight of the aircraft in multiple ways (such as both being able to alleviate flight loads and as well as have the ability to maneuver the aircraft.) In addition to continuous usage and morphing actuators built into the aircraft structure itself, continuous morphing technology also has the ability to actuate the vehicle across multiple different axes simultaneously. This can produce highly maneuverable aircraft, as well as more efficient aircraft.

Morphing aircraft technology also has potential benefits in an area of aircraft design known as “aero-servo-elasticity.” Aero-servo-elasticity is the tight coupling of aerodynamics, flight dynamics and control and aircraft structures. Due to their very nature, continuously morphing wing aircraft have the potential to allow for aircraft to be natively built exploit aero-servo-elastic synergies. For example, an active aero-servo-elastic system can be used to increase aircraft performance beyond a fixed (non-morphing design) through the use of much lighter

structures. Lighter, flexible structures have the tendency to undergo aerodynamic flutter, and aero-servo-elasticity can be used to prevent this. In addition, these aero-servo-elastic systems can be used for gust-load alleviation and/or high angle-of-attack maneuvering. This will allow the aircraft to be highly maneuverable, and very light (much like a bird.) It follows then that morphing wing technologies could allow for a new class of vehicles built around aero-servo-elasticity.

In addition, the use of flapping wing aircraft could also be investigated. All modern aircraft are built around the Cayley paradigm, which effectively decouples the aerodynamics and propulsion into two separate entities. However, the combined use of smart materials and novel actuator concepts could allow for efficient and silent flapping wing aircraft for the small scale (UAS). This could allow for stealthy, quiet and highly maneuverable aircraft, and as such, some effort should be focused on developing these aircraft.

However, the biggest technical challenge that must be overcome before continuously morphing aircraft become common place is in the technology of the actuator itself. This is discussed in the Structures section.

Multidisciplinary design optimization (MDO) is a field that couples computational aerodynamics, structures, propulsion and other related fields into one design problem. The main advantage of using MDO to design an aircraft system is that it exploits synergistic interactions between different subsystems to find an overall “globally optimized” design. The existence of global minima, brought through synergistic interactions can lead to extremely efficient new classes of aircraft. For example, NASA has investigated the coupled use of flow control and aerodynamics to create the X-57 SCEPTOR electric aircraft. NASA predicts that the X-57 should have about a 4.8x improvement over the state-of-the-art in energy usage during cruise compared to the state-of-the-art. These vast improvements would not be possible without the use of MDO.

While it seems as if the design of the aircraft using MDO clearly seems to be the future, there are many challenges that need to be overcome before it is possible to truly achieve a full optimization of the aircraft system. These problems can be broadly categorized into two separate focus areas of research: the development of MDO architectures and the development of high performance computer modeling. MDO architectures are schemes that couple the different building blocks or subsystems (such as aerodynamics, propulsion, etc.) into the global optimization scheme (the globally optimized aircraft design.) Many of the subsystems that are being modeled are very computationally expensive, and as such, can only fully run a handful of times given limited computer resources. These limited runs, thereby only generate limited data points. The challenge of an MDO architecture then becomes to use these sparse data points to try to identify the location of the global minima. There are many different types of MDO schemes, and all have varying degrees of accuracy, but more investigation into this will allow the designer to process through many types of design configurations and iterations effectively and more accurately while allowing for fewer calls to computationally expensive systems. For example, promising approaches to expanding the capabilities of MDO architectures may be through the use of machine learning.

The second challenge that needs to be solved before widespread usage of MDO is to develop and simplify the building blocks of the optimization – namely the aerodynamics, propulsion, structural and control systems solvers. Computational aerodynamics / fluid dynamics (CFD) is the primary bottleneck in MDO. It is computationally expensive to accurately solve for a full aircraft in flight. There is continual research in the development of high order (highly accurate) methods that can capture relevant physics within lower computational resources. Another challenge is the close coupling between subsystems. Often times, many of the schemes are interdependent on each other, and need to be solved simultaneously. For example, fluid-structure interaction, which forms the basics of aero-elasticity, requires the close coupled solving of both fluid and structural equations of motion. This can be very difficult, and better methods to efficiently link the solvers could be developed. In addition to software improvements, increased funding on high performance computing clusters (HPC) will allow for more complex MDO problems to be run, and will continue to allow the U.S. Air Force to maintain its technological edge over adversaries.

Lastly, other governmental agencies, such as NASA, are investing heavily into MDO research. In their recent report, NASA outlined that MDO, and other foundational building blocks would be a key focus area through the 2030 timeframe. Working closely with other governmental agencies to develop and maintain MDO frameworks should be prioritized.

Structures

Airframe structures of today's UAVs are mostly composed of composite materials such as carbon, aramid and glass fiber/polymer matrix as well as lightweight foam, plastic and lightweight metals. They are designed to safely carry the flight loads as well as being damage resistant during take-off, landing, transport, and storage. Wing, fuselage and tail-planes are all designed to be as light as possible so that the UAV can dedicate more of the total flight load capacity towards heavier and more sophisticated sensors and/or larger and higher capacity batteries or more fuel for longer missions. Additionally, for combat vehicles, airframe weight savings allow for more ammunition and or more sophisticated larger weapons that have higher lethality.

The idea of multifunctional airframe materials and structures is not new, but thus far the implementation of these technologies is very limited. Future research is required to enhance functionality by combining the structure with electronics, sensory functions, antennas, energy storage and actuators. This can be accomplished on the structural as well as material level. For instance on the structural scale, antenna wiring can be imbedded within the composite construction and routed in a pattern that is more efficient than traditional antenna configurations. On the material level the structure can be composed of load bearing material that is inherently capable of functioning as an antenna material. Likewise the structural material may also be capable of energy storage (made of battery material).

Integrated distributed sensors and actuators also will play an important role in future UAVs. Pitot-static air speed sensors are particularly vulnerable to damage and fouling. A typical pitot-static tube protrudes forward from the vehicle and is quite delicate and often is damaged during transport and storage. It can also be fouled by dust, water or other contaminants and thus provide erroneous airspeed measurements. The use of surface mounted/integrated air

speed sensing would be a major contribution to small UAS. Imbedded, solid-state actuators for distributed control could lead to a more weather resistant system than traditional servo actuated designs. Additionally, streamlining without hinge singularities leads to a more aerodynamics design and thus high efficiency.

The structural community has traditional emphasized damage resistance in the development of new materials. These materials have emphasized the improvement in specific materials properties such as tensile strength, compressive strength, and fracture toughness as an example. Although the goal of improving these properties and at the same time improving the specific strength by either decreasing density or improving the property itself will continue to be of high priority for aircraft the idea of damage tolerance is starting to become more relevant to aircraft. Damage tolerance allows for damage to transpire yet also allows for the structure to be viable after the damage has occurred. This can be accomplished by either routing the structural integrity around the damage site or healing the damage in-situ. Routing the loads around the site may mean that local damage results in less of a stress concentration as compared to traditional materials. Materials that can accomplish healing are currently being investigated and numerous mechanisms to accomplish this have been proposed. The technology has yet to be implemented on aircraft though. The panel suggests that more effort be put into self-healing materials.

Three-dimensional (3D) printing promises to significantly improve the ability to fabricate complex structural geometries that are difficult if not impossible to produce using other traditional manufacturing methods. This topic has relevance to UAVs as well as full-scale manned aircraft. It is covered in a previous section of this write-up. In the Mission Support section special treatment is given to 3D printing for “Field Repairs”.

The development of artificial muscle materials with high stroke, high strength and high bandwidth have alluded scientist and engineers to date. Most electromechanical actuators that are utilized in mechanical components derive their core kinematics from rotary actuation with transduction methods to produce linear motion. Engineers have been able to produce rotary based actuators with a wide range of force, stroke, and bandwidth. Using pneumatics linear motion is achievable that matches the needs of large scale aircraft, but does not scale well for small aircraft systems because of weight restrictions. Artificial muscle materials such as shape memory alloys, shape memory polymers, piezoelectric materials and electroactive polymers have been developed and implemented on small UAVs, but all suffer from performance deficiencies. Shape memory alloys have poor bandwidth and are slow, shape memory polymers are also slow and produce small forces, piezoelectric materials have high bandwidth and modest strength but have poor range of motion, typically far less than one percent elongation. Promises that were made, starting in the 1980s during the advent of the smart material revolution have not materialized for actuators. The promises in the sensory material however has been modestly realized. By combining actuator and sensor into one material, the full potential would be groundbreaking.

The development of the next generation of artificial muscle material will lead to a renaissance for UAVs as well as the entire field of robotics. They have the potential to replace common rotary based actuators, and hydraulic systems in traditional design and lead to fully integrated morphing wing vehicle. Research into artificial muscle materials must adhere to

specific and unique requirements associated with UAVs. Those requirements are high bandwidth (must be able to actuate in a fast manner), high actuation authority to weight ratio (adding weight to the UAV is not an acceptable trade-off), high deformation range (elongation/contraction should be greater than 5% full range), low power consumption, and low voltage (many of the actuators utilized today require more than 1,000 volts making them dangerous to handle and are incompatible with battery technologies), low levels of hysteresis (control of the vehicle is challenging enough without hysteresis).

There is an ever increasing need to expand the flight envelope of small UAVs. This can be accomplished using morphing where the UAV can change shape depending on the mission requirements. Shape changes may include wingspan, wing chord, wing sweep, wing airfoil as well as changes in the shape of the tail-plane and fuselage. These in-flight changes can be instantaneously accomplished to match the desired flight conditions. Such changes can be observed in nature where birds change their wing shapes for the different flight regimes including take-off-and-landing, cruise, rapid maneuvers and fast dash. Under the “Aerodynamics” heading of this section the benefits of morphing are discussed. Material and structural challenges are currently limiting factors for implementation. Smart material, discussed in the previous section hold the most promise in implementing morphing on UAVs. As a result future efforts are recommended in morphing aircraft using artificial muscle technology.

Materials utilized for UAVs must have high specific strength and stiffness. Since the 1970s with the advent of high modulus/high strength fiber reinforced carbon, aramid and glass fiber reinforced materials there have been only small improvements in materials used for aircraft. In the 1990’s and into this century nano-materials promised to drive the specific properties another order of magnitude higher. This promise is still not realized on the structural scale of the aircraft. Modern UAVs are incorporating more and more composite technology, but that technology was initially developed in the previous century and it is only now being utilized by improvements in fabrication technologies such as automated tape placement and related technologies.

Advancements in 3d printing are taking place at a high rate but the emphasis has concentrated on metals and plastic materials. These materials only represent a small fraction of those used for UAS. Therefore, it is recommended that an effort to develop 3D printing or a comparable rapid prototyping technology be developed using fiber reinforced or nano-reinforced materials. Additionally, the production of composite airframes is time consuming and requires expensive tooling and autoclave facilities. It is recommended that research to improve the fabrication of composite airframes be funded to combine rapid prototyping technology with autoclave free fabrication methodologies.

Propulsion

Propeller propulsion systems on current generation small unmanned air systems (sUAS) are optimized around a single flight condition. Naturally, operation of the aircraft outside of this design flight condition will be less efficient. For example, the propulsion system may be sized for optimal take-off behavior. Unfortunately, this has drawbacks for cruise flight, and the aircraft may be limited in range and endurance. Large aircraft are able to solve this problem by having propeller propulsion systems that can mechanically change their geometry to match an optimal

operating point for the flight condition. These adaptive systems are made possible through the use of variable pitch propellers. These systems are driven through pressurized oil or hydraulic systems that change the angle of the propeller blade as it rotates through the air. This ability for the propeller to change its pitch can allow the system to be optimized for different stages in flight – it can have a lower pitch during take-off and a higher pitch during cruise. Unfortunately, this kind of complexity just isn't practical for sUAS. This impracticality is largely due to the fragility and complexity of small mechanical systems. This drawback becomes a major concern as aircraft ruggedization is an implicit requirement for many sUAS that are being fielded today. However, if this problem is solved in a manner that is simple and effective, there would be completely new possibilities for sUAS aircraft. For example, it would be possible to have heavy lift hand-launchable sUAS that also can fly for much longer range and endurances compared to the state-of-the-art. In addition, the development of runway independent VTOL aircraft would also become much more efficient for long endurance missions as well.

Emerging technologies that might enable this capability are based in the development of smart materials, active flow control and novel control schemes. Recently, researchers were able to develop a swashplate-less UAV that was able to change the pitch of its propeller based on variable controller input. If this technology were given sufficient time to mature, it might become possible that this could be fielded in next generation sUAS. In addition, smart materials research could identify ways to change the shape of the propeller while maintaining mechanical simplicity. Active flow control could also be used to that effect, through “virtual” shaping of the propeller. All these technologies would allow for variable pitch propellers while maintaining mechanical simplicity.

Smaller micro air vehicles are envisioned in the future. These smaller vehicles could have the capability to infiltrate buildings and other tight spaces completely undetected. This could allow for precision strike capabilities within an urban environment or for covert reconnaissance. A major limiting factor for this type of capability is the miniaturization and complexity of the propulsion system. As discussed previously, the complexity requirements and size requirements are directly proportional. That is, as the size of the system reduces, so must its complexity, in order to maintain usefulness in the field. As such, there exists significant challenges in the size and practicality of propulsion systems if the current trend towards smaller vehicles is to be continued. Future solutions may include the use of flapping wing technology, flow control or smart materials.

One of the primary reasons of the rise in popularity of small unmanned air systems beginning in the late 1990s was due to the development of electric powertrains. These electric powertrains use battery technology as their main power source. As expected, these aircraft are also fundamentally limited in performance by these batteries. The main challenge this technology faces is the need to develop high energy density batteries while also maintaining sufficiently higher power density levels. Unfortunately, much work has been on the development of high energy density batteries, but there is often a trade-off as these batteries cannot sustain high levels of power output. Aircraft are notoriously high power consumers, so usually the trade-off favors high power density, with reasonable energy density for modern aircraft. This technology will need to mature if aircraft performance levels can approach that of their fossil fuel counterparts.

Fuel cells are devices which use the chemical reactions of fuel and oxidizer to generate electricity. This electricity then is either stored in a battery or is fed directly into the propulsion system. While fuel cells have been flown successfully on sUAS aircraft, they have not yet seen widespread adoption. Part of the problem is that they often use hydrogen, which is difficult to store and produce, as a fuel source. Next, the conversion process from hydrogen to electricity causes the fuel cell to become bulky and heavy. This extra size and weight penalty will lead to poor performance for the UAS. Sufficient technology research and maturation must still occur for this technology to see widespread adoption. If these challenges are overcome (fuel availability and storage, weight of the fuel cell), then ultra-long range UAS without having the need for high energy density batteries could be developed.

Lastly the widespread use of solar power could open new classes of long endurance UAS. Solar panels could harvest irradiated solar energy and store them on-board the UAS. If the solar energy harvested exceeds the energy consumed, then the aircraft could stay aloft indefinitely. The main challenge with this is that current generation solar panels are notoriously inefficient in their ability to absorb the solar irradiation (between 10 and 30%, on average.) With current levels of technology, it is difficult to have an indefinite UAS unless the aircraft is designed to have high levels of aerodynamic efficiency. This naturally leads to trade-offs in other performance metrics for the UAS. These trade-offs make solar power technology impractical for most UAS applications. However, if the efficiencies of these cells were to increase, then persistent UAS could be developed. More research should be focused on this area.

Controls

The topics on “Artificial Intelligence and Machine Learning, AI and Autonomy”, “Cyber Systems”, and “Sensing and Communication” can be applied to UAS. Therefore this section will only address the controls specific to the unmanned vehicle itself.

System identification is a cumbersome process used to fine-tune the inner loop of the control algorithms for aircraft. It is proposed that an effort be placed on the development of a robust real-time system ID methodology that can be implemented on a wide variety of vehicle configurations. The idea is for the autopilot to tune itself in-situ. The algorithm would account for changing conditions of the aircraft such as damage that occurs during flight, storage or takeoff and landing. The vehicle independent philosophy is consistent with the idea of a modular design that covers a wide ranging task (see mission support section). Additionally, the method would be more robust if it included airspeed sensor-less control, since airspeed sensors are notoriously vulnerable to damage.

The degree of autonomy of a UAS is typically bimodal, either being fully autonomous with very little input from the pilot during a mission, or completely controlled by a pilot with no preprogramming (tele-op mode). Having the ability to prescribe the degree of autonomy in real-time would be an important asset to the warfighter. The human-machine interface is a key feature of this technology and must be addressed. Additionally, having the ability to control the vehicle from a remote location, anywhere around the world, would benefit the warfighter by allowing for remotely located expertise to be involved in a mission.

Sense and avoid technology is key to the widespread reliance on UAVs. Ground based vehicles rely on LIDAR sensors, which are heavy and thus not applicable to a wide range of small UAS. Miniaturization of LIDAR would be extremely beneficial to UAS and would allow for a degree of autonomy not yet realized. Additionally, other miniature sensors that can be used for sense and avoid would be welcome.

Vehicle localization using GPS is the standard method utilized by modern UAS. The reliance on GPS has become an issue when GPS is denied, either through jamming or in environments where GPS reception is poor or nonexistent, such as under bridges or in buildings. Methodologies need to be developed for navigation and control in GPS denied environments.

Mission Support

A trend has emerged to have a wider variety of small unmanned systems that carry out a wider variety of mission requirements. Each has its own unique operating system, support equipment, training protocol, ground-station and storage requirements. This adversely affects implementation in the field and requires operators to be trained in an ever increasing regimen. For instance, vehicle requirements for long duration ISR are quite different than fast dash strike and thus today are performed by two completely different systems requiring training for both.

It is proposed that the range of mission and flight envelope be vastly increased using a single system (ground station, training, storage) using a modular design approach. This can be accomplished with a single versatile spine that can accommodate multiple wing configurations ranging from high speed to high endurance, multiple motor/propeller configurations, a variety of batteries, and multiple payloads that can include a variety of sensors for a variety of missions. The system would have a user interface that would have the warfighter enter the mission requirements (such as objective, flight path, range, and altitude) and the flight configuration would suggested, the flight planned in advance, and the recipe for the component combination to achieve the mission. The user would then quickly assemble the parts on the spine and launch the vehicle. This system could be used to replace the many small UAVs that operate within a size range that parallels that of the modular system. For instance the Puma, Raven, Switchblade, Wasp and Snipe would all be one modular system. Ultimately, this would require increased sophistication of the autopilot control algorithms, miniaturization of the avionics and advancements in many technologies such as materials, aerodynamics, control and propulsion.

3D Printed replacement parts are already being utilized for field repairs, but on an ad-hoc basis. If a small plastic part breaks in the field, which is a common occurrence during aircraft landings, the operator has two options; wait until a replacement part is delivered to the location, or perform a field repair. Getting parts to the arena is a dangerous and cumbersome process and sometimes not even possible. Field repairs often are only temporary and lack the fidelity that the aircraft requires to function as intended. It is proposed that an efficient process be developed to combine all part CAD/CAM files and 3D printing hardware in the field, along with directions for repair, and ready for implementation by the warfighter.

Storage and transport considerations are usually left as an afterthought for most UAVs, whereas an optimal system should include these considerations from the inception of the air vehicle concept. The preferred embodiment of a UAS would include provisions for compact

storage in a lightweight container. This would allow a fighting unit to carry UAS into the field more easily and thus allow for other provisions to be deployed or more UAS. This can be seen in many of today's UAS such as the snipe and the switchblade.

There are two predominate ways for a small UAV to remain unnoticed while in operation: disguising itself as something already in the environment (such as a bird) or blending into the background to appear invisible.

Biomimicry can be utilized not only to take cues from natural fliers to understand the aerodynamics of flight, but also to directly imitate the form and motion to appear just like one. Birds and bats are common throughout the world and most people don't pay much attention to indigenous species. If one could create a vehicle that looks and flies just like a common bird species it would decrease the likelihood of being noticed and thus could operate in hostile territory and be used as a surveillance platform more effectively. This is not a new idea, but one that has yet to be realized in practice. Vehicles that derive their propulsion from flapping are most bird-like and hence would be effective candidates. Flapping vehicles however have distinct audible signatures from the complex rotary transmissions and can be heard even before that can be seen. Noise signature of flapping vehicle would have to be addressed as well as form-factor. Large birds, such as vultures are ubiquitous around the world and typically spend most of their time soaring, without flapping. Additionally, they operate well about the ground and thus can be mimicked with a fixed wing vehicle with conventional propeller drive configuration. The flight pattern of a soaring bird can be more easily mimicked than that of a flapping bird.

Appendix I - Summit Details

The two-day event was hosted at the University of Florida (Gainesville, FL) in the John W. Reitz Union, a modern facility in the heart of the academic campus with state-of-the-art meeting facilities. The summit agenda follows:

APRIL 30, 2018			MAY 1, 2018		
7:30AM – 9:00AM	Registration & Check-In Continental Breakfast Networking & Poster Displays	Reitz Union Grand Ballroom	7:30AM – 8:00AM	Check-In Continental Breakfast Networking & Poster Displays	Reitz Union Grand Ballroom
9:00AM – 12:00PM	Plenary Session Introductory welcome, background perspective, and statement of summit goals. Video presentation by US Air Force Topic introductions by technology team facilitators	Reitz Union Grand Ballroom	8:00AM – 12:00PM	Parallel Panel Sessions Continuation of technical discussions, Report Out from Day 1, and refinement of technical recommendations and consensus views.	Reitz Union Various Locations
12:00PM – 1:00PM	Group Lunch View Ted Talk	Reitz Union Grand Ballroom	12:00PM – 1:00PM	Group Lunch Brief Recap	Reitz Union Grand Ballroom
1:00PM – 4:00PM	Parallel Panel Sessions In depth discussions clustered around broader themes, with dedicated subject matter experts, facilitators and reporters.	Reitz Union Various Locations	1:00PM – 4:00PM	Parallel Summary Sessions Each technical topic that has emerged will meet to finalize S&T forecasting and recommendations. Team leads will record and draft preliminary consensus statements.	Reitz Union Various Locations
4:30PM – 5:30PM	Plenary Wrap-Up Report out on technical progress, setting stage for Day 2	Reitz Union Grand Ballroom	4:30PM – 5:30PM	Plenary Wrap-Up Technical Progress	Reitz Union Grand Ballroom
6:00PM – 7:30PM	Social Event & Networking	UF Stadium Champions Club			

Appendix II - Summit Topics

- 1. Advanced modeling** - Mechanics, fluids and reacting flows, high stress- strain mechanics, constitutive modeling, structural coupling, damage models, theoretical and computational domains.
- 2. Advanced and Additive Manufacturing (AM)** - Vehicles, propulsion systems, rapid prototyping, design for AM, materials assembly, and custom assembly of matter.
- 3. Excited State Materials Performance** - Materials for customized and adaptive effects for increased materials performance, reduced cost and enhanced performance materials, reconfigurable and collaborative materials performance.
- 4. Energetics/Fuels** - Advanced propellants and energetic materials, including form, manufacture, and energy release, new materials, higher efficiencies, renewable fuels, high energy density fuels, energy storage, and fuel flexibility.
- 5. Hypersonics and High-speed Effects** - Fundamentals of high-speed vehicles and challenges, optimal control, next generation vehicles, manned and unmanned systems, command and control, propulsion, flow control, and aeroelasticity.
- 6. Cyber Systems** - Applied Cryptography and Privacy, Cloud and Distributed Systems Security, Internet and Mobile Security, Cyber Physical Systems Security, Systems and Storage Security, Internet-of-Things (IOT) Security, Supply Chain Security, System-on-Chip (SoC) Security and Trust, Biometrics and Security, Verification of IP Security and Trust Assured autonomous systems and networks, reliance to and detection of cyber effects/data corruption.
- 7. Machine Learning, Autonomy and AI** - Adaptive machine learning methods, Deep Learning, technologies to enable scaled autonomy, technologies for computational efficiency for embedded systems, formal methods for validation and verification of autonomous behavior.
- 8. Sensing and Communication** - Advanced, asynchronous, and intermittent feedback methods, control theory, scalability, human machine interaction, adaptation, learning, and artificial intelligence methods, quantification, autonomous operation in communication/sensing denied regions, Interoperability of heterogeneous networks, and remote sensing to improved situational awareness.
- 9. Space Systems** - Small satellites, space systems, space debris, orbital control, constellations, low-cost systems, and cyber effects.
- 10. UAVs** - Advanced concepts for propulsion, agility, and deployment.

Appendix III – List of Participants

The following individuals – listed in alphabetical order – participated in the discussions and/or contributed to the final report:

First Name	Last Name	Affiliation
Ossama	Abdelkhalik	Michigan Tech University
Robert	Andersen	55th ISS (RC-135/RIVET JOINT)
S.	Balachandar	UF
Mark	Balas	Univ of Tennessee
Prabir	Barooah	UF
Riccardo	Bevilacqua	UF
David	Bourell	UT Austin
Kevin	Bowcutt	Boeing Co.
Ashley	Carlton	MIT
Charles	Carpenter	Lockheed Martin
Anthony	Castrogiovanni	Acent Labs
Oana	Cazacu	University of Florida
Suman	Chakravorty	Texas A&M University
Youping	Chen	UF
John	Conklin	UFL
Antonio	de la Serna	Draper
Melanie	DeProspero	UF
Tarik	Dickens	FAMU-FSU College of Engineering/HPMI
Moses	Divaker	UF
Warren	Dixon	UF
Avinash	Dongare	University of Connecticut
Tarek	Elgohary	UCF
Jim	Fletcher	UNF
Joe	Foster	AF (Retired)
Chris	Goyne	University of Virginia
Hitomi	Greenslet	UF
David	Hahn	UF
Mohamed	Hashim	Resthetics
Jeffrey	Haynes	Aerojet Rocketdyne
Kristy	Hollingshead	IHMC
Ryan	Houim	UF
Diane	Howard	ERAU
Yong	Huang	UF
Peter	Ifju	UF
Greg	Ingram	Forge.AI
Thomas	Jackson	UF
Moriba	Jah	UT Austin

First Name	Last Name	Affiliation
David	Jeffcoat	UF REEF
Robert	Kielb	Acent Labs
Ted	Krauthammer	UFL
Ashok	Kumar	UF
Anastasios	Lyrintzis	Embry-Riddle Aeronautical University
Reda	Mankbadi	Embry Riddle Aeronautical University
Sean	Meyn	UFL ECE
Miroslav	Pajic	Duke University
Panos	Pardalos	UF Industrial Engineering
Simon	Phillpot	UF
Jose	Principe	UF
Subramanian	Ramakrishnan	FAMU-FSU College of Engineering
Anil	Rao	UF
Benoit	Revil-Baudard	UFL
Chris	Salthouse	UFL FLARE
Corin	Segal	UF
John	Shea	UF
Glen	Simser	Northrop Grumman
Douglas	Spearot	UFL
Ankit	Srivastava	Texas A&M University
Hedinn	Steingrimsson	Skaggreind EHF
Cameron	Stewart	UF
D. Scott	Stewart	Univ of Illinois at Urbana-Champaign
Robert	Suter	Carnegie Mellon University
Michael	Tonks	University of Florida
Venket	Tangirala	GE Research Center
Ufuk	Topcu	UT Austin
Trase	Travers	Advanced Core Concepts
Garritt	Tucker	Colorado School of Mines
Baba	Vemuri	UF
George	Voyiadjis	Louisiana State University
Girish	Wable	Jabil
Yan	Wang	Georgia Institute of Technology
Bin	Wei	GE
Markus	Wilde	Florida Institute of Technology
Xiaolin	Xu	UF
Michael	Zavlanos	Duke University
Mike	Zigan	Lockheed Martin

Appendix IV - List of References and Bibliography

- Abbeel, P., Schulman, J., Singh, S., & Silver, D. (2017). Tutorial workshop on Deep Reinforcement Learning. International Conference on Neural Information Processing Systems.
- Abdelkhalik, O. and Shadi Darani (2016). Hidden genes genetic algorithms for systems architecture optimization. In Proceedings of the Genetic and Evolutionary Computation Conference 2016, GECCO'16, pages 629–636, New York, NY, USA, ACM.
- Abdelkhalik, O. (2013). Hidden genes genetic optimization for variable-size design space problems. *Journal of Optimization Theory and Applications*, 156(2):450–468.
- Abdelkhalik, O. and Shadi Darani (2018). Evolving hidden genes in genetic algorithms for systems architecture optimization. *ASME Journal of Dynamic Systems, Measurement and Control*, in press.
- Abdelkhalik, O. (2011). Multi-gravity-assist trajectories optimization: Comparison between the hidden genes and the dynamic-size multiple populations genetic algorithms. In AAS/AIAA Astrodynamics Specialist Conference, number AAS11-620, Girdwood, Alaska, USA, July 31 – August 4, 2011.
- Abusoltan, M., S. Chakravorty and S. Khatri (2016). A GPU based implementation of Sensor Scheduling for Space Situational Awareness, FUSION 2016.
- Altermatt, E. R., & Pomerantz, E. M. (2005). The implications of having high-achieving versus low-achieving friends: A longitudinal analysis. *Social Development*, 14 (1), 61–81.
- Audley, H., Babak, S., Baker, J., Barausse, E., Bender, P., Berti, E., and Caprini, C. (2017). Laser Interferometer Space Antenna. arXiv preprint, arXiv:1702.00786.
- Bergman and A. Doucet (2000). Markov chain Monte Carlo data association for target tracking,” in Proc. ICASSP.
- Bertsekas, D., and Tsitsiklis., J. N. (1996). *Neuro-Dynamic Programming*. Cambridge, Mass: Atena Scientific.
- Blackman, S. S. and R. Popoli (1999). *Design and Analysis of Modern Tracking Systems*. Norwood, MA: Artech House.
- Blackwell, W. J., and Cahoy, K. (2017). New Small Satellite Capabilities for Microwave Atmospheric Remote Sensing. In *Seventh Conference on Transition of Research to Operations*.
- Blair, Y. and E. W. D. Blair (2000). *Multitarget Multisensor tracking: Applications and advances*, vol. III. Norwood, MA: Artech House.
- Borghesi, G. and Bellan, J. (2015). A priori and a posteriori investigations for developing large eddy simulations of multi-species turbulent mixing under high-pressure conditions,” *Physics of Fluids*, 27, pp. 035117.
- Borkar, V. S., and Meyn, S. P. (2012). Oja’s algorithm for graph clustering, Markov spectral decomposition, and risk sensitive control. *Automatica*, 48(10):2512–2519.
- Borkar, V. S., and Meyn., S. P. (2000). The ODE method for convergence of stochastic approximation and reinforcement learning. *SIAM J. Control Optim.*, 38(2):447–469.

- Brooks, J., and Barooah, P. (2016). Consumer-aware distributed demand-side contingency service in the power grid. *IEEE Transactions on Control of Network Systems*.
- Burr, J.R. and Yu, K.H (2017). Detonation propagation in a linear channel with discrete injectors and side relief,” 26th International Colloquium on the Dynamics and Explosion of Reactive Systems, July 30 – Aug. 4, 2017, Boston, MA.
- Cahoy, K. L., Douglas, E., Carlton, A., Clark, J., and Haughwout, C. (2017). How CubeSats contribute to Science and Technology in Astronomy and Astrophysics. In *American Astronomical Society Meeting Abstracts (Vol. 229)*.
- Chakraborty, R. A. (2015). An Efficient Recursive Algorithm for Atlas Construction. *Mathematical Foundations of Computational Anatomy* .
- Chakraborty, R., and Vemuri, B. C. (2015). Recursive Fréchet Mean Computation on the Grassmannian and Its Applications to Computer Vision. *ICCV*.
- Chakraborty, R., Seo, D., and Vemuri, B. C. (2016). An efficient Exact-PGA algorithm for constant curvature manifolds. In *Proceedings of the IEEE Conference of Computer Vision and Pattern Recognition*, pp. 3976--3984.
- Christy, R.W., and Agnar Pytte (1965). The Structure of Matter: Introduction to Modern Physics, Library of Congress 65-10938.
- Cheng, G., Salehian, H., and Vemuri, B. C. (2012). Efficient recursive algorithms for computing the mean diffusion tensor and applications to DTI segmentation. *ECCV*.
- Cong, S. et al. (2004). Markov chain Monte Carlo approach for association probability evaluation,” *Proc. IEE*, vol. 151, pp. 185–193.
- CPUOS (2011). Report of the Scientific and Technical Subcommittee on its forty-eighth session, held in Vienna from 7 to 18 February 2011, A/AC.105/987 at 48, available online at: http://www.unoosa.org/pdf/reports/ac105/AC105_987E.pdf (date last accessed: 27 April 2018).
- CPUOS (2017). Guidelines for the long-term sustainability of outer space activities: Working paper by the chair of the Working Group on the Long-term Sustainability of Outer Space Activities, A/AC.105.C.1/L.354 available online at: http://www.unoosa.org/res/oosadoc/data/documents/2017/aac_1052017c_11/aac_105c_11_354_0_html/AC105_C1_L354E.pdf (date last accessed: 27 April 2018).
- Dahm, W. (2010). A Vision for Air Force Science and Technology 2010–30.
- Darani, S. and Ossama Abdelkhalik (2016). Developments on the optimization of interplanetary trajectories using hidden genes genetic algorithms. In *Proceedings of The AIAA/AAS Astrodynamics Specialist Conference*, Long Beach, California, 13–16 September 2016. American Institute of Aeronautics and Astronautics (AIAA).
- Darani, S. and Ossama Abdelkhalik (2018). Convergence analysis of hidden genes genetic algorithms in space trajectory optimization. *AIAA Journal of Aerospace Information Systems*, in press.
- Darani, S. and Ossama Abdelkhalik (2018). Space trajectory optimization using hidden genes genetic algorithms. *Journal of Spacecraft and Rockets*, in press.

- Deptula, P., Rosenfeld, J., Kamalapurkar, R., and Dixon, W. E. (2018). Approximate Dynamic Programming: Combining Regional and Local State Following Approximations. *IEEE Trans. Neural Netw. Learn. Syst.*, 2154-2166.
- Devraj, A. M., and Meyn, S. P. (2016). Differential TD learning for value function approximation. In *55th Conference on Decision and Control*, pp. 6347–6354.
- Devraj, A. M., and Meyn, S. P. (2017). Fastest convergence for Q-learning. *ArXiv*.
- Devraj, A. M., and Meyn, S. P. (2017). Zap Q-learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*.
- Douglas, E. S., Bendek, E., Marinar, A., Belikov, R., Merck, J., and Cahoy, K. L. (2017). The DeMi CubeSat: Wavefront Control with a MEMS Deformable Mirror in Space. In *American Astronomical Society Meeting Abstracts (Vol. 229)*.
- Faber, W., S. Chakravorty and I. I. Hussein (2015). A Randomized Approach to the Solution of Multi-Target Tracking Problems, *FUSION 2015*.
- Faber, W., S. Chakravorty and I. I. Hussein (2015). A Comparison of RFISST and MHT for Space Situational Awareness Problems, *Proceedings 2015 AAS/AIAA Astrodynamics Specialists Conference*, Vail, CO.
- Faber, W., S. Chakravorty and I. I. Hussein (2016). An Improved Randomized Technique for Multi-Target Tracking, *FUSION 2016*.
- Faber, W., S. Chakravorty and I. I. Hussein (2018). Multi-Target Tracking using a Randomized Hypothesis Generation Technique, *IEEE Transactions on Aerospace and Electronic Systems*, under review.
- Fréchet, M. (1948). Les éléments aléatoires de nature quelconque dans un espace distancié. *Annales de l'institut Henri Poincaré* (pp. vol(10),pp.215-310). Presses universitaires de France.
- Friedkin, N. E., and Cook, K. S. (1990). Peer group influence. *Sociological Methods and Research*, 19 (1), 122-143.
- Fu, J., and Topcu., U. (2014). Probably approximately correct MDP learning and control with temporal logic constraints. *Conference on Robotics Science and Systems*.
- Fuster, J.M. (1995). *Memory in the Cerebral Cortex*. MIT Press.
- Gad, A. and Ossama Abdelkhalik (2011). Hidden genes genetic algorithm for multi-gravity-assist trajectories optimization. *AIAA Journal of Spacecraft and Rockets*, 48(4):629–641, July–August.
- Goodman, I.R., R. P. S. Mahler, and H. T. Nguyen (1997). *Mathematics of Data Fusion*. Kluwer Academic Publishers.
- Halevy A., N. P. (2009). The unreasonable effectiveness of data. *IEEE Int. Sys. Magazine*, 8-12.
- Hesamoddin Salehian, R. C. (2015). An efficient recursive estimator of the Fréchet mean on hypersphere with applications to Medical Image Analysis. *Mathematical Foundations of Computational Anatomy*.
- Ho, J., Cheng, G., Salehian, H., and Vemuri, B. (2013). Recursive Karcher Expectation Estimators And Geometric Law of Large Numbers. *AISTATS*, pp. 325-332.

- Houim, R.W. and Kuo, K.K. (2013). A ghost fluid method for compressible reacting flows with phase change,” *Journal of Computational Physics*, 235, pp. 865-900.
- Huang, W. C. (2011). Feature selection for neuro-dynamic programming. In F. Lewis, *Reinforcement Learning and Approximate Dynamic Programming for Feedback Control*. Wiley.
- Hussein, I., K. DeMars, C. Frueh, R. S. Erwin, and M. Jah (2012). An Aegis-fisst algorithm for joint detection and tracking in space situational awareness, *AAS/AIAA Astrodynamics Specialist Conference*, August.
- Izzo, D. (2017). Problem description for the 9th Global Trajectory Optimization Competition, May 2017.
- Kamalapurkar, R., Andrews, L., Walters, P., and Dixon, W. E. (2017). Model-based reinforcement learning for infinite-horizon approximate optimal tracking. *IEEE Trans. Neural Netw. Learn. Syst.*, 753-758.
- Kamalapurkar, R., Walters, P. S., Rosenfeld, J. A., and Dixon, W. E. (2018). Reinforcement learning for optimal feedback control: A Lyapunov-based approach. *Springer*.
- Khalil, A. E.E and Gupta, A.K. (2013). Fuel flexible distributed combustion for efficient and clean gas turbine engines,” *Applied Energy*, 109, pp. 267-274.
- Kingsbury, R. W., Nguyen, T., Riesing, K., and Cahoy, K. (2017). Fast-steering solutions for cubesat-scale optical communications. In *International Conference on Space Optics—ICSO 2014* (Vol. 10563, p. 105630G). International Society for Optics and Photonics.
- Kingsbury, R. W., Riesing, K. M., Cahoy, K. L., Nguyen, T. N. T., and Caplan, D. O. (2017). U.S. Patent No. 9,813,151. Washington, DC: U.S. Patent and Trademark Office.
- Knuth, J., and Barooah, P. (2013). Collaborative localization with heterogeneous inter-robot measurements by Riemannian optimization. *IEEE Conference on Robotics and Automation*, pp. 1534-1539.
- Knuth, J., and Barooah, P. (2014). Distributed Collaborative 3D Pose Estimation of Robots from Heterogeneous Relative Measurements: Optimization on Manifold Approach. *Robotica*, 1-29.
- Lanier, J. (2014). *Who Owns the Future?* Simon & Schuster.
- Le Bon, G. (1896). *The Crowd: A Study of the Popular Mind*. Viking Press.
- Mach, E. (2010). *The Science of Mechanics*, Fourth ed.; Watchmaker Publishing.
- Mahler, R. P. S. (2007). *Statistical Multisource-Multitarget Information Fusion*. Artec House.
- Mahler, R. P. S. (1999). Multitarget Markov motion models, *SPIE Conference on Signal Processing, Sensor Fusion, and Target Recognition*, vol. 3720, pp. 47–56.
- Marlow, W. A., Carlton, A. K., Yoon, H., Clark, J. R., Haughwout, C. A., Cahoy, K. L., and Morzinski, K. M. (2017). Laser-Guide-Star Satellite for Ground-Based Adaptive Optics Imaging of Geosynchronous Satellites. *Journal of Spacecraft and Rockets*, 1-18
- Mehta, P. G., and Meyn, S. P. (2009). Q-learning and Pontryagin’s minimum principle. In *IEEE Conference on Decision and Control*, pp. 3598–3605.
- Meyn, S. P. (2007). *Control Techniques for Complex Networks*. Cambridge University Press.

Mitchell, M., Hollingshead, K., and Coppersmith, G. (2015). Quantifying the language of schizophrenia in social media. Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality, pp. 11-20. Denver, CO.

National Science Foundation. Cyber Physical Systems Program.
http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=503286.

Nguyen, V.-A., Boyd-Graber, J., Resnik, P., and Miler, K. (2015). Tea party in the house: A hierarchical ideal point topic model and its application to Republican legislators in the 112th Congress. Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (ACL-IJCNLP), (pp. 1438–1448).

Oh, S. et al. (2009). Markov chain Monte Carlo data association for multi-target tracking,” IEEE Tr. Automatic Control, vol. 54, pp. 481–497.

Oran, E.S. (2015). Understanding explosions – From catastrophic accidents to creation of the universe,” Proceedings of the Combustion Institute, 35, pp. 1-35.

Pardalos, P., Conca, P., Giuffrida, G., and Nicosia, G. (2015). Machine Learning, Optimization and Big Data: First International Workshop.

Park, R. E., and Turner, R. H. (1967). On social control and collective behavior: selected papers. In The sociological writings of Robert E. Park, pp. 269-274. Chicago: University of Chicago Press.

Pauling, L. (1947). General Chemistry; Dover Pub (1969).

Pasula, H. et al. (1999). Tracking many objects with many sensors,” in Proc. IJCAI.

Plank, M. (1923). Treatise on Thermodynamics; Dover Pub (1969).

Principe, J., and Chalasani, R. (2014). Cognitive Architecture for Sensory Processing. Proceedings of the IEEE , 102 (4), 514-525.

Pryor, O., Barak, S., Koroglu, B. Ninnemann, E. Vasu, S.S. (2017). Measurement and interpretation of shock tube ignition delay times in highly CO₂ diluted mixtures using multiple diagnostics,” Combustion and Flame 180, pp. 63-76.

Radhakrishnan, A. D. (2016). Learning techniques for feedback particle filter design. In 55th Conference on Decision and Control, pp. 5453–5459.

Raihan, D. and S. Chakravorty (2015). A Hybrid Particle Filter-UKF algorithm for tracking Space Objects”, Proceedings 2015 AAS/AIAA Astrodynamics Specialists Conference, Vail, CO.

Raihan, D. and S. Chakravorty (2018). Particle Gaussian Mixture Filters I”. Automatica, in press.

Raihan, D. and S. Chakravorty (2018). Particle Gaussian Mixture Filters II”. Automatica, in press.

Reid, D.B. (1979). An algorithm for tracking multiple targets,” IEEE Transactions on Automatic Control, vol. 24, pp. 843–854.

Ristic, B. et al. (2010). Improved SMC implementation of the PhD filter,” in Proc. Conf. Info. Fusion (FUSION).

- Salehian, H., Chakraborty, R., Ofori, E., and Vemuri, B. (2015). An efficient recursive estimator of the Frechet mean on a hypersphere with applications to Medical Image Analysis. MFCA.
- Salehian, H., Cheng, G., Vemuri, B. C., and Ho, J. (2013). Recursive Estimation of the Stein Center of SPD Matrices and Its Applications. ICCV, pp. 1793--1800.
- Salehian, H., Vaillancourt, D., and Vemuri, B. C. (2014). iPGA: Incremental Principal Geodesic Analysis with Applications to Movement Disorder Classification. MICCAI, pp. 765-772.
- Serra, P., Barnwell, N., Ritz, T., and Conklin, J. (2017). Integrated, Software-Defined Pulse Modulator and Laser System for Small Satellite Optical Communication.
- Shalom, Y. B. and X. Li (1995). Multitarget-Multisensor Tracking: Principles and Techniques, New York.
- Shalom, Y. B. and T. Fortmann (1988). Tracking and Data Association. San Diego, CA: Academic Press.
- Skuhersky, M., Larissa Balestrero Machado, Markus Wilde, and Christopher Brett (2017). "CERBERUS: Prototype for an Agile Inspection and Servicing Satellite Using Thrust-Vectoring Cold-Gas Propulsion", AIAA SPACE and Astronautics Forum and Exposition, AIAA SPACE Forum, AIAA 2017-5116.
- Slater, J.C. (1939). Introduction to Chemical Physics; McGraw-Hill.
- Slater, J.C. (1960). Quantum Theory of Atomic Structure, McGraw-Hill.
- Strakey, P. (2018). Oxy-Combustion Flame Fundamentals for Supercritical CO₂ Power Cycles," The 6th International Supercritical CO₂ Power Cycles Symposium, March 27-29, 2018, Pittsburgh, PA.
- Stage, C., Dezechache, G., Pierre, J., and Grèzes, J. (2015). Emotional contagion: Its scope and limits. Trends in Cognitive Sciences , 19 (6), 297–299.
- Sunberg, Z., S. Chakravorty and R. Erwin (2013). Information Space Receding Horizon Control", IEEE Transactions on SMC-B, vol. 43, pp. 2255-2260.
- Sunberg, Z., S. Chakravorty and R. Erwin (2012). Multi-Agent Information Space Receding Horizon Control, Proceedings 2012 American Control Conference, Montreal, Canada.
- Sunberg, Z., S. Chakravorty and R. S. Erwin (2014). Information Space Receding Horizon Control for Space Situational Awareness", Proceedings of the American Control Conference, Portland, OR.
- Taghvaei, A., and Mehta., P. G. (2016). Gain function approximation in the feedback particle filter. In IEEE Conference on Decision and Control, pp. 5446–5452.
- Taghvaei, A., Mehta, P. G., and Meyn., S. P. (2017). Error estimates for the kernel gain function approximation in the feedback particle filter. In Proceeding of the American Control Conference.
- Taheri, E., Di Wu, Chi-En Lee, Jared Shimoun, Ossama Abdelkhalik, Shangyan Zou, Shadi Darani, Brandon Jackson, and Jacob Liimatta (2017). Gtoc9: Results from team Michigan Tech University of Michigan.

- Taylor, B.D., Kessler, D.A. Gamezo, V.N., and Oran, E.S. (2014). Numerical simulations of hydrogen detonations with detailed chemical kinetics,” Proceedings of the Combustion Institute, 34, pp. 2009-2016.
- Thampan, T., Shah, D., Cook, C., Novoa, J., and Shah, S. (2014). Development and Evaluation of Portable and Wearable Fuel Cells for Soldier Use.” J. Power Sources Volume 259, 1 August, pp. 276-281.
- Turner, R., & Killian, L. M. (1987). Collective Behavior. Prentice Hall.
- Vo, B. and W. Ma (2006). The Gaussian mixture probability hypothesis density filter, IEEE Transactions On Signal Processing, vol. 54, no. 11, November.
- Vo, B. et al. (2005). Sequential Monte Carlo methods for multi target filtering with random finite sets, IEEE Tr. Aerosp. Electronic Systems, vol. 4, pp. 1224–1245.
- Vo, B.T. et al. (2007). An analytic implementation of the Cardinalized probability hypothesis density filter, IEEE Tr. Signal Processing, vol. 55, pp. 3553–3567.
- Wen, M., and Topcu, U. (2016). Probably approximately correct learning in stochastic games with temporal logic specifications. In Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence.
- Wen, M., Papusha, I., and Topcu, U. (2017). Learning from demonstrations with high-level side information. Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence.
- Weyl, H. (1919). The Continuum – A Critical Examination of the Foundations of Analysis, Dover Pub (1994).
- Wilde, M., Ian Walker, Stephen Kwok Choon, and James Near (2017). "Using Tentacle Robots for Capturing Non-Cooperative Space Debris - A Proof of Concept", AIAA SPACE and Astronautics Forum and Exposition, AIAA SPACE Forum, AIAA 2017-5246.
- Wolanski, P. (2013). Detonative Propulsion, Proceedings of the Combustion Institute, 34, pp. 125-158.
- Yang, V. (2000). Modeling of supercritical vaporization, mixing, and combustion processes in liquid-fueled propulsion systems, Proceedings of the Combustion Institute, 28, pp. 925-942.