



**NAVAL
POSTGRADUATE
SCHOOL**

MONTEREY, CALIFORNIA

COMMAND AND CONTROL MODELS OF NEXT GENERATION

UNMANNED AIRCRAFT SYSTEMS

by

W. David Place & Dr. Mark E. Nissen

October 2015

Approved for public release; distribution is unlimited

Prepared for: The Consortium for Robotics and Unmanned Systems Education and Research (CRUSER)

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 12-10-2015		2. REPORT TYPE Technical Report		3. DATES COVERED (From-To) 10/01/2014 – 12/31/2015	
4. TITLE AND SUBTITLE Command and Control Models of Next Generation Unmanned Aircraft Systems				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) W. David Place & Dr. Mark E. Nissen				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943				8. PERFORMING ORGANIZATION REPORT NUMBER NPS-IS-15-002	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) The Consortium for Robotics and Unmanned Systems Education and Research (CRUSER)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The technologic capabilities of autonomous systems (AS) continue to accelerate, and integrated performance by AS and people working together can be superior to that of either AS or people working alone. We refer to this increasingly important phenomenon as Teams of Autonomous Systems and People (TASP), and through our recent research—representing the current state of the art—we have demonstrated computational experimentation capability in the TASP domain. The problem is, several technology trends suggest that unmanned aircraft may be diverging away from operating and behaving like their manned counterparts, suggesting that some of our most futuristic model assumptions may be off target. This is where our ongoing research project begins to make an important contribution by investigating the next generation of autonomous systems. In this technical report, we motivate and introduce such TASP research, and we provide an overview of the computational environment used to experiment on TASP command and control organizations and phenomena. We summarize in turn the research method. Key results follow, and we conclude then by summarizing our agenda for continued research along these lines.					
15. SUBJECT TERMS Artificial intelligence; autonomy; command & control; experimentation; modeling & simulation; unmanned aircraft systems.					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Mark Nissen
a. REPORT	b. ABSTRACT	c. THIS PAGE			
Unclassified	Unclassified	Unclassified	UU	65	831-656-3570

THIS PAGE INTENTIONALLY LEFT BLANK

**NAVAL POSTGRADUATE SCHOOL
Monterey, California 93943-5000**

Ronald A. Route
President

Douglas A. Hensler
Provost

The report entitled “*Command and Control Models of Next Generation Unmanned Aircraft Systems*” was prepared for and funded by The Consortium for Robotics and Unmanned Systems Education and Research (CRUSER).

Further distribution of all or part of this report is authorized.

This report was prepared by:

Dr. Mark E. Nissen
Professor

W. David Place
Research Associate

Reviewed by:

Released by:

Dan Boger, Chairman
Information Sciences

Jeffrey D. Paduan
Dean of Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

The technologic capabilities of autonomous systems (AS) continue to accelerate, and integrated performance by AS and people working together can be superior to that of either AS or people working alone. We refer to this increasingly important phenomenon as Teams of Autonomous Systems and People (TASP), and through our recent research—representing the current state of the art—we have demonstrated computational experimentation capability in the TASP domain. The problem is, several technology trends suggest that unmanned aircraft may be diverging away from operating and behaving like their manned counterparts, suggesting that some of our most futuristic model assumptions may be off target. This is where our ongoing research project begins to make an important contribution by investigating the next generation of autonomous systems. In this technical report, we motivate and introduce such TASP research, and we provide an overview of the computational environment used to experiment on TASP command and control organizations and phenomena. We summarize in turn the research method. Key results follow, and we conclude then by summarizing our agenda for continued research along these lines.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I. INTRODUCTION.....	1
A. AUTONOMOUS SYSTEMS	1
B. OPEN C2 QUESTIONS	2
C. COMPUTATIONAL EXPERIMENTATION.....	3
D. RESEARCH OVERVIEW.....	5
II. BACKGROUND	7
A. POWER COMPUTATIONAL ENVIRONMENT	7
B. POWER MODEL EXAMPLE	11
C. KEY RESULTS.....	16
1. Autonomy Degree.....	16
2. Interdependence Level.....	17
3. C2 Implications	17
III. RESEARCH METHOD	21
A. METHOD SUMMARY	21
IV. RESULTS	23
A. RESULTS OVERVIEW.....	23
B. TECHNOLOGY	23
1. Artificial Intelligence	24
2. Machine Learning.....	25
3. Biologic Design	28
4. Computational Hardware	30
5. Perception	31
C. PROGRAMMATICS	34
D. EDUCATION AND TRAINING	37
E. CULTURE.....	38
F. OPERATIONS	39
G. STRATEGY.....	41
V. CONCLUSION	45
VI. REFERENCES.....	49
INITIAL DISTRIBUTION LIST	55

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

A. AUTONOMOUS SYSTEMS

The US Department of Defense (DoD), along with the militaries of NATO members and other allied nations, has discovered and begun to capitalize upon the value of robots, unmanned vehicles and other autonomous systems (AS) for a variety of different missions, ranging from search and rescue, through aerial bombing, to Cyberspace surveillance. To a large extent, people in such military organizations operate and control the AS, much the same way that people in many factories operate and control machines for production, assembly and packaging. The AS are basically slaves to their human operators.

The technologic capabilities of AS continue to accelerate, however, and systems in some domains have reached the technical point of total autonomy: they can perform entire missions without human intervention or control. For instance, in 2001 a Global Hawk flew autonomously on a non-stop mission from California to Australia, making history by being the first pilotless aircraft to cross the Pacific Ocean (AMoD, 2001). As another instance, in 2013 a Northrop Grumman X-47B unmanned combat air vehicle successfully took off from and landed on an aircraft carrier underway at sea (BBC, 2013).

This elucidates many emerging issues in terms of command and control (C2). Who, for instance, commands and controls unmanned aircraft when they fly autonomously? Clearly there are operators who monitor such vehicles, and there are commanders who authorize their missions, but the mission itself is conducted autonomously, and it remains somewhat unclear whom to hold accountable (e.g., the commander, the operator, the engineer, the manufacturer) if something goes wrong or whom to credit if all goes well.

Further, as technologic sophistication continues to advance rapidly (e.g., in computational processing, collective sense making, intelligent decision making), a wide array of diverse robots (e.g., in hospitals; see Feil-Seifer & Mataric, 2005), unmanned vehicles (e.g., for highway driving; see Muller, 2012) and other intelligent systems (e.g., for industrial control; see McFarlane et al., 2003) continue to demonstrate unprecedented capabilities for extended, independent and even collective decision making and action

(e.g., offensive and defensive swarming; see Bamberger et al., 2006). Indeed, the technologic maturity of many AS available today (e.g., UCLASS – Unmanned Carrier-Launched Airborne Surveillance and Strike; see Dolgin et al., 1999) exceed the authority delegated to them by organizations and leaders; that is, their performance is limited more by policy than technology (e.g., see DoDD 3000.09, 2012).

In many skilled mission domains and under demanding environmental conditions (e.g., tactical surveillance; see Joyce, 2013), AS are replacing people at an increasing rate (e.g., unmanned vs. manned aircraft sorties; see Coutts, 2012). These machines can outperform their human counterparts in many dimensions (e.g., consistency, memory, processing power, endurance; see Condon et al., 2013), yet they fall short in other ways (e.g., adaptability, innovation, judgment under uncertainty; see HRW, 2012). Task performance by AS is optimal in some situations, and performance by people is best in others, but in either case, the respective capabilities of autonomous machines and people remain complementary. As such, *integrated* performance, by complementary autonomous systems and people *working together*, can be superior in an increasing number of circumstances, including those requiring skillful collective action (Nissen & Place, 2013).

Hence there is more to this trend than simple technologic automation of skilled work by machines (e.g., numerical control machining) or employment of computer tools by skilled people (e.g., computer aided drafting). Where autonomous systems and people collaborate together in coherent teams and organizations, we refer to this increasingly important phenomenon as Teams of Autonomous Systems and People (TASP).

B. OPEN C2 QUESTIONS

TASP raises a plethora of open, C2 research, policy and decision making questions. For one, under what circumstances should people work subordinate to AS (e.g., robot supervisor) versus controlling them (e.g., robot subordinate)? Few researchers, policy makers or organization leaders are even asking this question today, much less trying to answer it, as the conventional, conservative and often naïve bias is overwhelmingly toward people controlling machines. Nonetheless, empirical evidence shows that AS can produce superior results—in some circumstances—when people are subordinate (e.g., see Bourne, 2013). This represents revolutionary change, and our

millennia of accumulated knowledge in terms of C2, organization, management, leadership, information science, computer science, human-systems integration and like domains leaves us largely unprepared to seize upon such situated performance superiority.

For another, under what circumstances should units comprised of people be organized, led and managed separately from counterparts comprised of AS (e.g., separate aircraft squadrons), and what circumstances favor instead organization integration¹ of people and AS into combined units (e.g., integrated or composite squadrons; see CFFC, 2014)? Because every mission-environment context manifests some uniqueness, the answer may vary across diverse missions, environments, times and organizations; even individual personnel skills, team trust levels, leadership characteristics, political risk aversion, and like factors may affect the approach leading to greatest mission efficacy. Indeed, a central aspect of mission planning and execution may require explicit consideration of how people and AS should be organized, and such TASP organization may even require dynamic replanning and change mid-mission.

For a third, how can researchers, policy makers and leaders develop confidence that their chosen C2 organization approach (e.g., to subordinating or superordinating robots to people, to separating or integrating AS and personnel units, to selecting missions involving collaboration between people and AS) will be superior? These technology-induced research questions are so new and foreign that negligible theory is available for guidance, and it is prohibitively time-consuming, expensive and error-prone to systematically test the myriad different approaches via operational organizations. This is the case in particular where loss of life, limb or liberty may be at stake.

C. COMPUTATIONAL EXPERIMENTATION

Computational experimentation offers an unmatched yet largely unexplored potential to address C2 questions along these lines. If computational models can be

¹ For instance, HSM-35, located at NAS North Island, has been organized and configured to manage and support both the Fire Scout UAS and the H-60 aircraft (e.g., integrated technicians and operators have been trained to maintain and operate both systems). Additional information and guidance is available in the USFF/CNAF UAS Concept of Operations. Nonetheless, several questions remain: Is such integration a good idea? On what science is it based? What are the comparative advantages and disadvantages? How could it become even more effective?

developed to represent the most important aspects of organizations with existing, planned or possible TASP benefits, then researchers could employ such models to address the kinds of open questions posed above. Moreover, organization leaders, managers and policy makers could develop confidence in their situated decisions and actions involving the organization, integration and leadership of AS and people.

Further, once such computational models have been developed and validated, they can become virtual prototype C2 organizations to be examined empirically and under controlled conditions through efficient computational experiments (e.g., see Oh et al., 2009). Indeed, tens, hundreds, even thousands of diverse approaches to TASP C2 can be examined very quickly, with their relative behavior and performance characteristics compared to match the best C2 approach with a variety of different missions, environmental conditions, technologic capabilities, autonomy policies, personnel characteristics, skill levels and job types. Moreover, such computational experimentation and comparison can be accomplished very quickly and at extremely low cost relative to that required to experiment with teams or organizations in the laboratory—or especially in the field—with no risk of losing life, equipment or territory in the process (e.g., see Nissen & Buettner, 2004).

Toward this end, recent research (Nissen & Place, 2014)—representing the current state of the art (i.e., VDT; see Levitt et al., 1999)—has employed computational modeling and simulation technology to demonstrate computational experimentation capability in the TASP domain. Specifically, an agent-based model, which captures and reflects the structure, behavior and performance characteristics of C2 organizations in the field, is used to examine alternate C2 approaches and AS capabilities—both as available today and projected for the future—in the context of exploring TASP opportunities, alternatives and decision spaces.

Computational models assess six distinct degrees of AS capability—ranging from *Degree 0 – no autonomy* (i.e., manned aircraft) to *Degree 5 – future capability* (e.g., AS matching manned capabilities)—corresponding to both current and potential ship and aircraft platforms (i.e., CVN, DDG, LCS, F/A-18, MH-60, Scan Eagle, Fire Scout, Triton, future AS). Models also assess four distinct levels of mission interdependence—ranging from *Pooled* (e.g., manned *or* unmanned missions in separate airspaces) to

Integrated (e.g., manned *and* unmanned missions in common airspaces)—as an orthogonal dimension. For instance, consider the nature of a mission—and its corresponding C2 requirements and complications—with an **unmanned wingman** flying alongside a manned aircraft leader, or symmetrically, consider a manned pilot flying as wingman alongside his or her unmanned aircraft leader.

Together these six degrees of autonomy and four levels of interdependence produce a 24 cell matrix of TASP contexts that are assessed in terms of C2 organization performance for a 24-hour ISR mission. Each cell is represented by a separate computational model, which is simulated 50 times, across eight performance dimensions, to create a substantial analytic space.

However, such recent research has assumed that next generation unmanned aircraft will operate and behave increasingly like manned aircraft. The problem is, several technology trends suggest that unmanned aircraft may be diverging instead of converging, developing their own, unique modes of operation and sets of behaviors. Indeed, some such modes and sets may make integration of manned and unmanned aircraft more challenging, not less. This has enormous C2 implications, and our recently developed computational modeling and simulation capability hinges on the results of the current research project.

D. RESEARCH OVERVIEW

This is where our ongoing research project continues to make an important contribution. The project described in this technical report centers on expanding our recently enabled C2 computational modeling and simulation capability to understand next generation unmanned aircraft systems, with particular emphasis on specifying advanced models for computational experimentation. In particular, we investigate technology trajectories and design visions for next generation unmanned aircraft systems through qualitative methods. Specific techniques include archival review, semi-structured interviews and participant observation. We also leverage results to specify corresponding computational models for extended experimentation in the TASP domain. This involves two primary tasks: 1) investigate the technology trajectories and design visions for next generation unmanned aircraft systems; and 2) understand how to represent such

trajectories and visions in terms of agent-based computational models. To scope the project appropriately, we concentrate on Task 1 in this investigation.

In the balance of this technical report, we first provide an overview of the POWer computational experimentation environment and summarize results from our experiments on TASP C2 organizations and phenomena, for these constitute key background for the current study. Then we summarize the research method. Key results follow, and we conclude in turn with our agenda for continued research along these lines.

II. BACKGROUND

A. POWER COMPUTATIONAL ENVIRONMENT

This section draws heavily from Gateau and colleagues (2007) to provide an overview of the POWER computational environment. POWER builds upon the planned accumulation of collaborative research over roughly two decades to develop rich, theory-based models of organization processes (Levitt, 2004). Using an agent-based representation (Cohen, 1992; Kunz et al., 1999), micro-level organization behaviors have been researched and formalized to reflect well-accepted organization theory (Levitt et al., 1999). Extensive empirical validation projects (e.g., Christiansen, 1993; Thomsen, 1998) have demonstrated the representational fidelity and shown how the qualitative and quantitative behaviors of our computational models correspond closely with a diversity of enterprise processes in practice.

This research stream continues today with the goal of developing new micro-organization theory and embedding it in software tools that can be used to design organizations in the same way that engineers design bridges, semiconductors or airplanes—through computational modeling, analysis and evaluation of multiple virtual prototypes. Such virtual prototypes also enable us to take great strides beyond relying upon the kinds of informal and ambiguous, natural-language descriptions that comprise the bulk of organization theory and C2 doctrine today.

For instance, in addition to providing textual description, organization theory is imbued with a rich, time-tested collection of micro-theories that lend themselves to computational representation and analysis. Examples include Galbraith's (1977) information processing abstraction, March and Simon's (1958) bounded rationality assumption, and Thompson's (1967) task interdependence contingencies. Drawing on such micro-theory, we employ symbolic (i.e., non-numeric) representation and reasoning techniques from established research on artificial intelligence to develop computational models of theoretical phenomena. Once formalized through a computational model, the symbolic representation is “executable,” meaning it can be used to emulate organization dynamics.

Even though the representation has qualitative elements (e.g., lacking the precision offered by numerical models), through commitment to computational modeling, it becomes semi-formal (e.g., most people viewing the model can agree on what it describes), reliable (e.g., the same sets of organization conditions and environmental factors generate the same sets of behaviors) and explicit (e.g., much ambiguity inherent in natural language is obviated). This, particularly when used *in conjunction with* the descriptive natural language theory of our extant literature, represents a substantial advance in the field of organization analysis and design, and it offers direct application to research and practice associated with C2.

Additionally, when modeling aggregations of people—such as work groups, departments or firms—one can augment the kind of symbolic model from above with certain aspects of numerical representation. For instance, the distribution of skill levels in an organization can be approximated—in aggregate—by a Bell Curve; the probability of a given task incurring exceptions and requiring rework can be specified—organization wide—by a distribution; and the irregular attention of a worker to any particular activity or event (e.g., new work task or communication) can be modeled—stochastically—to approximate collective behavior. As another instance, specific organization behaviors can be simulated hundreds of times—such as through Monte Carlo techniques—to gain insight into which results are common and expected versus rare and exceptional.

Of course, applying numerical simulation techniques to organizations is hardly new (Law and Kelton, 1991), but this approach enables us to *integrate* the kinds of dynamic, qualitative behaviors emulated by symbolic models with quantitative metrics generated through discrete-event simulation. It is through such integration of qualitative and quantitative models—bolstered by reliance upon sound theory and empirical validation—that our approach diverges most from extant research methods and offers new insight into organization and C2 dynamics.

We summarize the key POWER elements via Table 1 for reference. Most of these elements are discussed below, but this table provides a concise summary. The interested reader can refer to the work by Gateau and colleagues (2007) for details.

Table 1 POWer Elements and Descriptions

Model Element	Element Description
Tasks	Abstract representations of any work that consumes time, is required for project completion and can generate exceptions.
Actors	A person or a group of persons who perform work and process information.
Exceptions	Simulated situations where an actor needs additional information, requires a decision from a supervisor, or discovers an error that needs correcting.
Milestones	Points in a project where major business objectives are accomplished, but such markers neither represent tasks nor entail effort.
Successor links	Define an order in which tasks and milestones occur in a model, but they do not constrain these events to occur in a strict sequence. Tasks can also occur in parallel. POWer offers three types of successor links: finish-start, start-start and finish-finish.
Rework links	Similar to successor links because they connect one task (called the <i>driver</i> task) with another (called the <i>dependent</i> task). However, rework links also indicate that the dependent task depends on the success of the driver task, and that the project's success is also in some way dependent on this. If the driver fails, some rework time is added to all dependent tasks linked to the driver task by rework links. The volume of rework is then associated with the project error probability settings.
Task assignments	Show which actors are responsible for completing direct and indirect work resulting from a task.
Supervision links	Show which actors supervise which subordinates. In POWer, the supervision structure (also called the exception-handling hierarchy) represents a hierarchy of positions, defining who a subordinate would go to for information or to report an exception.

POWer has been developed directly from Galbraith’s information processing view of organizations. This view of organizations, described in detail by Jin and Levitt (1996), has three key implications.

The first is ontological: we model knowledge work through interactions of *tasks* to be performed; *actors* communicating with one another and performing tasks; and an *organization structure* that defines actors’ roles and constrains their behaviors. Figure 1 illustrates this view of tasks, actors and organization structure. As suggested by the figure, we model the organization structure as a network of reporting relations, which can capture micro-behaviors such as managerial attention, span of control and empowerment. We represent the task structure as a separate network of activities, which can capture organization attributes such as expected duration, complexity and required skills. Within the organization structure, we further model various *roles* (e.g., marketing analyst, design engineer, manager), which can capture organization attributes such as skills possessed,

levels of experience and task familiarity. Within the task structure, we further model various sequencing constraints, interdependencies and quality/rework loops, which can capture considerable variety in terms of how knowledge work is organized and performed.

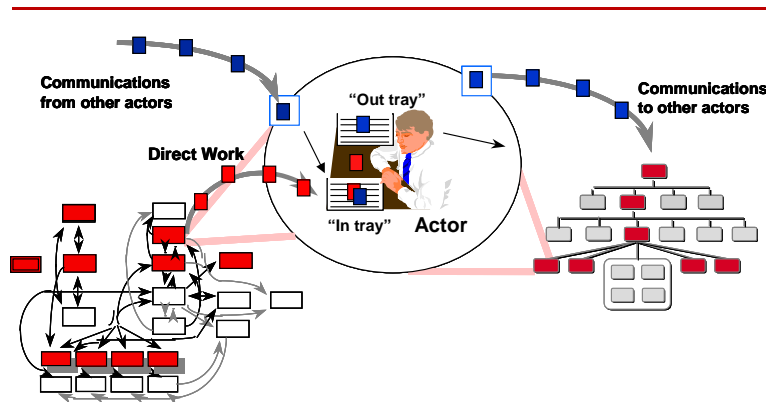


Figure 1 Information Processing View of Knowledge Work

As suggested by the figure also, each actor within the intertwined organization and task structures has a queue of information tasks to be performed (e.g., assigned work activities, messages from other actors, meetings to attend) and a queue of information outputs (e.g., completed work products, communications to other actors, requests for assistance). Each actor processes such tasks according to how well the actor's skill set matches those required for a given activity, the relative priority of the task, the actor's work backlog (i.e., queue length), and how many interruptions divert the actor's attention from the task at hand.

The second implication is computational: *work volume* is modeled in terms of both *direct work* (e.g., planning, design, manufacturing) and *indirect work* (e.g., decision wait time, rework, coordination work). Measuring indirect work enables the quantitative assessment of (virtual) process performance (e.g., through schedule growth, cost growth, quality).

The third implication is validational: the computational modeling environment has been validated extensively, over a period spanning more than two decades, by a team of over 30 researchers (Levitt 2004). This validation process has involved three primary

streams of effort: 1) internal validation against micro-social science research findings and against observed micro-behaviors in real-world organizations, 2) external validation against the predictions of macro-theory and against the observed macro-experience of real-world organizations, and 3) model cross-docking experiments against the predictions of other computational models with the same input data sets (Levitt et al., 2005). As such, ours is one of the few, implemented, computational organization modeling environments that has been subjected to such a thorough, multi-method trajectory of validation.

B. POWER MODEL EXAMPLE

To help set up an example from our previous research project, here we summarize the experiment design, which centers on the two key dimensions: *Autonomy* and *Interdependence*. Briefly as outlined above, on the autonomy dimension we account for the technologic sophistication of the UAVs (Degree 0 – 5); on the interdependence dimension we account for the interdependence between multiple aircraft in concurrent operation (pooled, sequential, reciprocal, integrated), including both manned-only, unmanned-only and integrated manned-unmanned missions. With these two dimensions, we consider—in a systematic and orderly manner—a 6x4 matrix of increasingly complex TASP C2 contexts, which comprise collectively our set of experiment conditions.

We summarize this context matrix in Table 2. At the one extreme, we consider two manned aircraft that are deployed in separate geographical regions of controlled airspace (e.g., within the vicinity of its host ship) or in the same geographical region but at different times. This corresponds to Degree 0 autonomy with pooled interdependence (i.e., labeled “D0P” in the table). At the other extreme, we consider a squadron of completely autonomous UAVs and a squadron of manned aircraft flying integrated missions in uncontrolled airspace. This corresponds to a group of Degree 5 UAVs reflecting both reciprocal interdependence among themselves and integrated interdependence with their manned aircraft counterparts (i.e., labeled “D5I” in the table). Each of the key intermediate conditions (i.e., Degree 0 to Degree 5 autonomy, across all four interdependence conditions) is examined systematically also for completeness. This matrix summarizes our computational experiment design.

Table 2 TASP C2 Computational Experiment Design Summary

Degree\Interdependence	Pooled	Sequential	Reciprocal	Integrated
Degree 0	D0P	D0S	D0R	D0I
Degree 1	D1P	D1S	D1R	D1I
Degree 2	D2P	D2S	D2R	D2I
Degree 3	D3P	D3S	D3R	D3I
Degree 4	D4P	D4S	D4R	D4I
Degree 5	D5P	D5S	D5R	D5I

Each of these 24 experiment design cells is represented by a separate computational model, which is simulated 50 times, across eight performance dimensions (i.e., *duration, rework, coordination, wait, work cost, functional risk, mission risk and maximum backlog*), to create a substantial analytic space. In this previous study, we examine each of these 24 test cases in terms of *extant C2* organizations and approaches. In follow-on work, we plan to examine each of these cases in terms of the *best C2* organizations and approaches.

Figure 2 delineates a screenshot of our baseline CTG organization and platform set. The light (green) person icons represent organizations at four levels (i.e., CTF, CTG, Platform [e.g., DDG, LCS], Aircraft Operators [e.g., F/A-18, MH-60]). The dark (brown) rectangle icons represent operational leadership, decision making and staff work in addition to common tasks (e.g., planning, maintenance, air traffic control), whereas the light (yellow) rectangle icons represent the aircraft ISR mission tasks; each aircraft must take off, navigate to its area of interest, operate in ISR mode, and then return to the ship for landing or recovery. Organizations and tasks are represented at appropriate levels: sufficiently low to capture the important structural and behavioral dynamics, but sufficiently high to abstract away details that do not impact the results in terms of TASP C2.

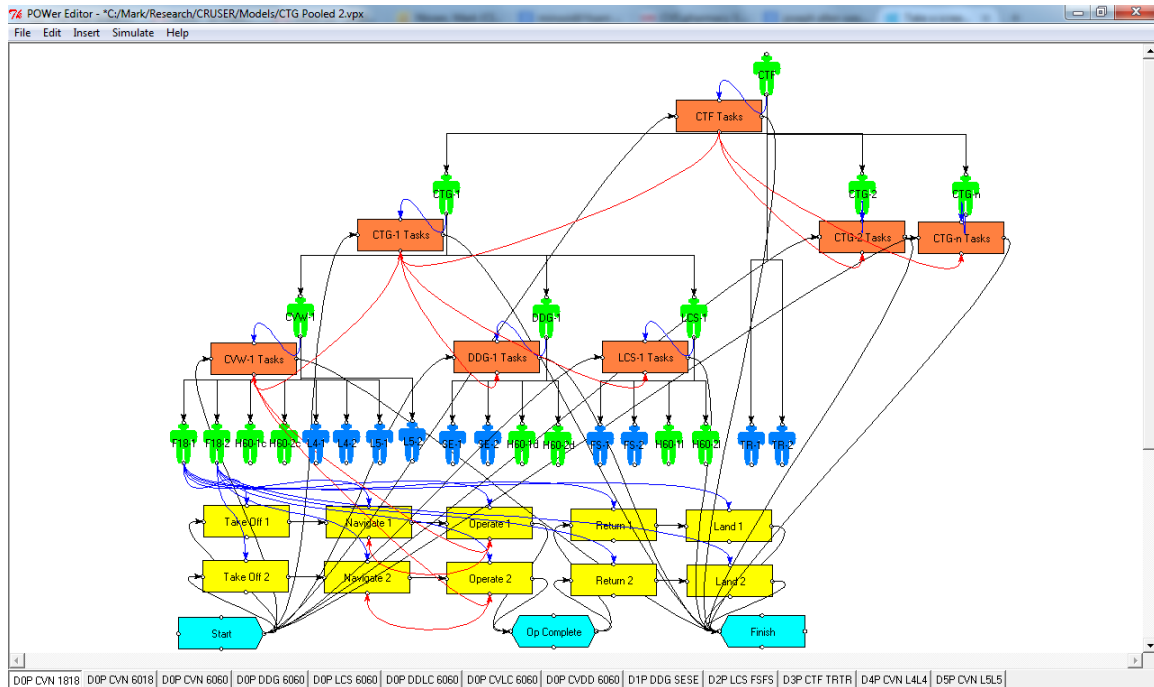


Figure 2 Baseline Model (DOP) – CVN F/A-18s

At the lowest level in the organization lies an array of diverse manned (light/green) and unmanned (dark/blue) aircraft. F/A-18s (Degree 0) are assigned to the CVN. MH-60s (Degree 0) are assigned to one or more DDGs and LCSs as well as the CVN. ScanEagles (Degree 1) are assigned to the DDGs, and Fire Scouts (Degree 2) are assigned to the LCSs. Tritons (Degree 3) are examined as an asset from beyond the CTG itself (e.g., controlled by the CTF), and we examine two future AS (Degree 4 & 5 UAVs) principally in terms of CVN assignment here². The many lines linking various icons in the figure are used to symbolize organization hierarchy, job assignment, task precedence, communication and other important model relations. For instance, the (dark/red) links connecting the Operate and Navigate tasks denote *rework*; if the Operate task fails to produce satisfactory ISR results (e.g., a promising contact is not located, insufficient intelligence is gathered, sensor data cannot be relayed), then the aircraft may have to

² Understanding that the corresponding Degree 4 and 5 technology has yet to be fielded and developed, respectively, these UAVs could be either fixed or rotary wing (or both), and hence could potentially operate effectively from the CVN and other ship platforms (esp. DDG, LCS).

Navigate to some other region and Operate there. The interested reader can refer to Gateau et al. (2007) for detailed explanations for all key model links and parameters.

In the baseline screenshot above, two (manned) F/A-18s are assigned to fly ISR missions in separate airspaces (i.e., D0P: Degree 0 autonomy, pooled interdependence). This task assignment is evident from the five (dark/blue) links between each F/A-18 actor and the aircraft ISR mission tasks (e.g., Take Off, Navigate, Operate); the first actor (labeled “F18-1” in the figure) is assigned to the upper sequence of tasks (e.g., labeled “Take Off 1,” “Navigate 1,” “Operate 1”), and the second actor (labeled “F18-2” in the figure) is assigned to the lower sequence of tasks (e.g., labeled “Take Off 2,” “Navigate 2,” “Operate 2”). Here both (manned) aircraft are assigned to the same (CVN) platform and (CVW) organization, and each flies in a different region of airspace (pooled interdependence). This represents a very common and relatively straightforward C2 context.

As noted above, the simulated ISR mission has a *planned* duration³ set at 24 hours. For aircraft such as the F/A-18s depicted in this model, such nominal 24 hour duration exceeds the endurance of a single aircraft sortie, so a sequence of sorties must be planned to span the whole 24 hour period, and sorties may have to continue beyond 24 hours in order to accomplish all mission objectives. We take into account each aircraft’s performance characteristics (esp. endurance) when specifying the computational model, and we record each aircraft’s simulated performance level (e.g., actual mission duration) in the computational experiment.

The other experiment conditions are specified and modeled similarly. In total we develop, specify, execute and analyze more than 36 different TASP C2 models to cover the 24 experiment conditions. Table 3 summarizes this model set. The first column lists each of the six autonomy degrees (e.g., “D0” = Degree 0; “D5” = Degree 5), and the second column designates which ship platforms⁴ the ISR aircraft operate from (e.g., “CVN” = carrier, “DDG” = destroyer, “LCS” = littoral combat ship, “LCDD” = littoral

³ Not all missions are equally effective, however, so some may take less than 24 hours to accomplish all ISR objectives successfully, whereas other may require (much) more time to complete.

⁴ The Triton is land-based, and we presume that it comes under CTF control, hence the “CTF” designation.

combat ship + destroyer⁵). The third column designates which two aircraft types conduct each ISR mission (e.g., “1818” = two F/A-18s, “6060” = two MH-60s, “60SE” = one MH-60 and one ScanEagle).

The remaining three columns are marked (x) to indicate where computational models have been developed. In all models where two aircraft of the same autonomy degree conduct a mission (e.g., two F/A-18s, two MH-60s, two ScanEagles), we examine all three interdependence levels (i.e., pooled, reciprocal, integrated⁶), but where a combination of manned and unmanned aircraft conduct a mission together (e.g., one MH-60 and one ScanEagle or Fire Scout, one F/A-18 and one Triton, one MH-60 or F/A-18 and one Level 4 or 5 UAV⁷), by definition only the integrated interdependence level applies, and hence only a single model is developed.

⁵ Where two different ship platforms are involved, we use only two letters for each, and we begin with the ship associated with the Leader aircraft. For instance, “LCDD” signifies that the Leader aircraft flies from an LCS and that the Wingman aircraft flies from a destroyer, whereas “DDLCS” indicates that the Leader aircraft flies from a destroyer and that the Wingman aircraft flies from an LCS.

⁶ Speaking technically, integrated interdependence does not apply to missions flown solely by two manned *or* two unmanned aircraft; that is, integrated interdependence applies only to missions flown by both manned *and* unmanned aircraft. Nonetheless, several model parameter settings differ between reciprocal and integrated interdependence experiment conditions, and it is informative to examine even all-manned or all-unmanned aircraft missions through both such conditions.

⁷ An implicit assumption is that the Level 4 UAV will be rotary wing helo, and hence tend to fly with the MH-60, and that the Level 5 UAV will be fixed wing jet, and hence tend to fly with the F/A-18. Nonetheless, *for C2 purposes*, we also examine some different combinations (e.g., Level 4 or 5 could be a vertical take-off and landing jet).

Table 3 Summary of TASP C2 Models

Level	Ship	Model Summary			
		Aircraft	Pool	Recip	Integ
D0	CVN	1818	x	x	x
	CVN	6060	x	x	x
	DDG	6060	x	x	x
	LCS	6060	x	x	x
D1	DDG	SESE	x	x	x
	DDG	60SE			x
	LCDD	60SE			x
D2	LCS	FSFS	x	x	x
	LCS	60FS			x
	DDLC	60FS			x
D3	CTF	TRTR	x	x	x
	CVCT	18TR			x
D4	CVN	L4L4	x	x	x
	CVN	60L4			x
	DDCV	60L4			x
D5	CVN	L5L5	x	x	x
	CVN	18L5			x
	CVDD	18L5			x

C. KEY RESULTS

Here we summarize key results of our C2 computational experimentation, which we organize in three parts: 1) Autonomy Degree, 2) Interdependence Level, and 3) C2 Implications.

1. Autonomy Degree

Regarding autonomy degree, we can generalize to say that none of our performance measures varies linearly with increasing autonomy. Nonetheless, autonomy degree is a very impactful variable, and we find major C2 performance differences across the various manned and unmanned aircraft types, particularly in terms of the measures *duration*, *functional risk* and *work cost*. Specifically, the UAVs in our current inventory (Autonomy 1 – 3) take considerably longer to complete the ISR mission than either the manned (Autonomy 0) or future unmanned (Autonomy 4 & 5) aircraft. These same, current inventory UAVs also present greater functional risk. Both of these results stem from the lesser skill and experience levels associated with the corresponding UAV

aircrews, with an accumulation of mistakes and required rework affecting mission duration and functional risk levels directly. This affects C2 substantially, as longer durations and greater functional risk levels must be accommodated explicitly through mission planning, execution, monitoring and intervention activities.

Work cost results affect C2 substantially also, as different aircraft can incur dramatically different operating costs for performing the same, nominal 24 hour ISR mission. Succinctly, the F/A-18 appears to represent the most costly ISR platform, yet it also exhibits the greatest speed and mission flexibility (e.g., strike). The MH-60 costs roughly half as much to operate, is comparable to the Fire Scout in that regard, and has distinct capabilities (e.g., rescue operations). The Triton costs in turn about half as much as the Fire Scout, and the ScanEagle—along with the Level 4 & 5 UAVs—is expected to cost even much less to operate. All other considerations aside, the operating cost of an ISR mission could become an important C2 consideration.

2. Interdependence Level

Regarding interdependence level, we isolate and examine interdependence effects, which are relatively clear: greater levels of interdependence correspond to higher values (but not necessarily “better” or “worse”) across almost all performance measures. This is clearly understandable in our TASP C2 context, for interactions between aircraft increase in frequency and intensity, and having multiple aircraft operating together in common airspace complicates their planning, operating, tracking, monitoring and intervening.

For instance, we find that missions take longer to complete with increasing interdependence levels. This is evident in particular with the transition from reciprocal to integrated interdependence. As another instance, we see clearly an even more extreme effect in terms of coordination. Coordinating multiple aircraft flying in common airspace has a major impact in terms of C2. A third instance centers on mission risk, which nearly doubles as we transition from reciprocal to integrated interdependence. C2 for integrated interdependence appears to have critical issues.

3. C2 Implications

Regarding C2 implications, we begin by summarizing findings pertaining to current aircraft in operational use today: there is little cause for concern regarding Autonomy 0 (e.g., F/A-18, MH-60) with pooled interdependence (e.g., operating in

separate airspaces). This is something that the Navy knows how to do well. There are no C2 implications per se here. This is business as usual.

Nonetheless, even though C2 appears to function acceptably well at present, several aspects of the extant and ubiquitous, military C2 organization and approach suggest that problems will emerge with continued advances in and integration of AS technology. For instance, the C2 organization reflects a tall, functional hierarchy, with considerable centralization, substantial formalization and frequent staff rotation. This makes for relatively long information flows and decision chains, coupled with perennial battles against knowledge loss from personnel turnover and challenges with cross-functional (and even more so with joint and coalition) interaction.

As another instance, the formalization inherent within this C2 organization reflects a strong dependence upon written standards, rules and procedures, but the pace at which UAVs and other AS are being introduced and integrated appears to be accelerating, and such formalization through written documents may have a hard time keeping up with rapid and local changes onboard various ships and among diverse aircrews. Many organization experts would argue that both instances militate against efficient—or even effective—C2.

Indeed, these study results imply—somewhat counter intuitively—that unmanned (ISR) missions require *more* planning, monitoring, intervening and like control activities than their manned counterparts. Given our current C2 organization and approach, greater numbers of C2 staff—or more skilled and experienced staff members—will be required for unmanned than for manned missions, and such missions will take more time, suffer from more mistakes, and generally tax the C2 organization more. Although the capabilities of future UAVs may mitigate these effects to some extent, we must continue to address and endure the higher C2 load for now, and we should consider redesigning our familiar, military C2 organization to address the imminent shortcomings noted above.

Further, required C2 skills and experiences may not be uniform across manned and unmanned missions. Indeed, we may find one set of C2 personnel (e.g., planners, controllers, watch standers) that is proficient principally in terms of manned missions, whereas a different cadre of personnel becomes proficient principally in terms of unmanned missions. *Given our current C2 organization and approach*, this will limit the

degree of flexibility available in terms of assigning *suitably experienced personnel* to different jobs, and many organizations will need to staff themselves with seemingly redundant personnel: some possessing skill and experience with manned operations and others possessing similar yet distinct skill and experience with their unmanned counterparts. Results reveal that as task interdependence continues its shift toward integrated manned-unmanned missions, such similar yet distinct skill and experience will likely break down and become ineffective for C2. Alternatively, as noted below, other approaches to organizing and conducting C2 offer potential to mitigate these detrimental effects.

Results reveal further that C2 organizations will make more mistakes; experience increasing time pressure; require greater effort, more time and higher cost to conduct missions; and operate under conditions of substantially higher mission risk. Moreover, if our C2 organization and approach appear fragile with aircraft operating reciprocally (or integrally) only in pairs today, then one can easily imagine how such organization and approach could break under the load of tens or even hundreds (or possibly thousands) of aircraft operating simultaneously and reciprocally (or integrally). Imagine further the exacerbation stemming from a shift to strike and other missions that diverge from the ISR context of this study.

Organizations will need to learn how to learn more quickly, and the current approach to education and training will likely fail. Other, more advanced approaches to accelerating knowledge flows through C2 organizations will likely become mandatory—so that people, teams and organizations can learn more quickly and with fewer mistakes—and people will need ways to learn just in time (JIT), knowing what to do and how to do it well in the local context (e.g., when and where such awareness is needed). Our extant military C2 organization and approach appear to be unprepared to meet these organization learning demands.

We understand increasingly well how C2 should focus on four, fundamental and inextricable elements: *people, process, organization* and *technology*. Our current C2 organization and approach operate well across many conditions and circumstances at present, but teams of autonomous systems and people—especially as manifested through

integrated manned-unmanned aircraft missions—appear to fall way beyond this set of current conditions and circumstances.

We do not have all of the answers in the previous study, but C2 organization (re)design lies at the center—particularly to promote rapid organization learning and to accelerate knowledge flows—and we may experience the compelling need to shift away from our familiar, hierarchical, archetypical C2 organization and toward higher maturity, agile, edge-like archetypes (e.g., *Collaborative*; see Alberts & Nissen, 2009).

Moreover, this elucidates an organization challenge for TASP C2 in general and the CTG in particular: mission efficacy may require shifting from one C2 approach and organization to another depending upon the context; that is, the same CTG (or other C2 organization) may need to employ different C2 approaches and organizations across the range of diverse TASP contexts, even within the same ISR mission set. Some nominal mission (e.g., “Mission-1”) may be approached best by the traditional hierarchy, for instance, but then the next nominal mission (e.g., “Mission-2”) may require Self-Synchronization or other, different C2 organization and approach. This will require not only understanding *which* C2 organization and approach is most appropriate for each particular mission, but also knowing *how* to transition from one C2 organization and approach to another. As such *we’re defining the state of the art with our research*, and such C2 organization selection and transition is *way beyond current practice*.

We need to understand more about future AS, and we need to examine current and future, manned and unmanned, aircraft missions through alternate C2 organizations and approaches (e.g., Hierarchy, Collaborative, Self-Synchronization). This represents the research trajectory on which this present study falls, as we examine future capabilities in considerable depth.

III. RESEARCH METHOD

A. METHOD SUMMARY

As noted above, we investigate technology trajectories and design visions for next generation unmanned aircraft systems through qualitative methods. Specific techniques include archival review, semi-structured interviews and participant observation. In terms of archival review, we focus on the current literature regarding the development of UAS with embedded artificial intelligence (AI) capabilities. Our experience in this field has led us to target a variety of sources, which serve to inform the study well and guide our qualitative investigation.

Semi-structured interviews target systems design engineers and organizations who are leading the efforts in this field. Our experience in this field has led us to focus on several organizations that are advancing the state of the art regarding UAS. We identify the key people within each organization and arrange to conduct both semiformal and informal interviews with each of them. Because many aspects of advanced UAS technology are classified, interviews are not transcribed, but we take thorough notes, and we summarize our thoughts rigorously following each interview session. Our interview frame is summarized in Table 4.

Table 4 Interview Frame

Role	Organization
Associate Research Professor	Carnegie Mellon University
Director, CyLab Mobility Research Center	Carnegie Mellon University
Senior Research Engineer	Stanford University
Distinguished Engineer	IBM Watson Group
Director, Government Solutions	IBM Watson Group
Director, Waterloo Autonomous Vehicles Laboratory (WAVELab)	University of Waterloo
Interoperability Systems Engineer	OSD
Director, Mobile Robot Laboratory	Georgia Institute of Technology
Systems Engineer	Naval Air Warfare Center
Systems Engineer	Naval Air Systems Command
Program Manager	DARPA Defense Sciences Office

Participant observation involves participation in numerous autonomous systems conferences. We're able to attend information sessions; converse with speakers, panelists and other attendees; and leverage such opportunities to identify additional sources of literature to review and to arrange additional interviews. Our conference frame is summarized in Table 5. For reference, conference organizers include Technology Training Corporation (TTC), Association of Unmanned Vehicles Systems International (AUVSI), US Office of Secretary of Defense (OSD), New Mexico State University (NMSU), US Navy Pacific Fleet (PACFLT) and the Massachusetts Institute of Technology (MIT).

Table 5 Conference Frame

Organization	Conference
TTC	Unmanned Aircraft Systems
TTC	Unmanned Aircraft Systems Payloads
TTC	Intelligence, Surveillance and Reconnaissance
AUVSI	unmanned systems Program Review
AUVSI	unmanned systems North America
OSD/NMSU	Tactical Analysis and Assessment Center
PACFLT	Autonomous Off-board unmanned systems Workshop
AAAI	Association for the Advancement of Artificial Intelligence
MIT	San Diego Executive Forum

Qualitative data are collected, coded and analyzed continuously, with the data collection effort not stopping until theoretical saturation is reached. Coding adheres to well-accepted grounded theory building techniques (Glaser & Strauss, 1967; Strauss & Corbin, 1998), including open, axial and selective coding activities in series.

IV. RESULTS

A. RESULTS OVERVIEW

In this section we present the key findings and results of our qualitative investigation. We organize this discussion into the six primary themes that emerge from our coding and analysis of qualitative data: 1) technology, 2) programmatic, 3) education and training, 4) culture, 5) operations, and 6) strategy. We integrate key quotations and examples where appropriate.

B. TECHNOLOGY

In this section we elaborate our analyses and insights pertaining to technology. From our readings, interviews and observations, it becomes apparent that the U.S. military's investments in unmanned aircraft systems (UAS) have proven invaluable for missions from intelligence, surveillance and reconnaissance (ISR) to tactical strike. Most of the current systems, however, require constant control by a dedicated pilot and sensor operator as well as a large number of analysts, all via telemetry⁸. These requirements severely limit the scalability and cost-effectiveness of UAS operations and pose operational challenges in dynamic, long-distance engagements with highly mobile targets in contested electromagnetic environments.

When considering the role that emerging technology will play in our future Naval Forces, it is clear that "Unmanned" Vehicles have a large and growing impact. This is reinforced by the 2009 CNO Strategic Studies Group (SSG) report (Hogg, 2009). We are rapidly approaching the point where every naval vessel will employ one or more UxVs⁹, and their employment as organic assets to individual fielded warfighters will become the norm. The term *unmanned* is a misnomer at present, however. A more appropriate term would be *uninhabited*, because, although there is no human physically in or on the vehicle, people are an integral component of all unmanned systems. In fact, the number of humans required to operate and exploit the capabilities of these vehicles is a major area of concern. So the question is not *if* humans are part of the system, but *what is their*

⁸ Very few systems have the capability to dynamically modify their flight profiles based on rudimentary logic algorithms, however none of these systems are based purely upon Artificial Intelligence.

⁹ The term UxV is often used to refer to all classes of unmanned vehicles (air, surface, underwater, ground).

role. This leads to the concept *level of autonomous behavior*, which we address through the study in terms of our dimension *autonomy degree*.

We identify five major, complementary, interrelated drivers of increasing autonomy. These include artificial intelligence (AI), machine learning, biologic design, computational hardware, and perception. For ease of organization and exposition, we address each in turn.

1. Artificial Intelligence

AI has been the focus of active research for many decades, but it appears to have made some major breakthroughs in the past few years. In artificial intelligence, the term autonomy implies bounded independent thought and action. As a fundamental principle, Simon's Law of Bounded Rationality 11 (Simon, 1991) states that the actions of a program or robot are bounded by the information it has, the amount of time available for computation, and the limitations of its algorithms. Thus, the independence of an AS is fixed by the designers.

AI refers to the common core of programming principles as "agency." If it is necessary to identify that an algorithm is restricted to a particular type of agent, AI refers to a mobile robot as a "physically-situated agent" to distinguish it from a "software agent," and "robot" is reserved for a system using the factory automation style of programming. (Defense Science Board, 2012)

One example of a mature robotic technology is the TurtleBot (Clear Path Robotics, 2009). These are commercially available off-the-shelf products. They come with a suite of sensors including but not limited to a Kinect sensor, eye view, gyroscope, odometry sensors on its wheels, and it drives like a car. On top of the TurtleBot there is a laptop that contains and applies the planning, control and other algorithms. The TurtleBot can operate autonomously. For several instances, it uses its Kinect sensor to get depth information about the walls and surroundings; it senses where walls and other obstacles are located; it calculates its physical location on an internal map representation; and it moves, collision-free, on its own: the only guidance needed is an ultimate destination.

2. Machine Learning

Machine learning can be viewed as a branch of AI with explicit emphasis on how AS acquire the knowledge needed for autonomous operation. Traditionally knowledge acquisition has been a long and laborious process, through which operators, analysts, engineers and programmers hand code the algorithms and enter ‘chunks’ of knowledge manually. This process becomes arduous for large or complex systems, and it remains infeasible in circumstances where subject matter experts (SMEs) are unwilling or unable to codify their knowledge in terms that can be processed by machines.

An exciting alternate approach is to produce computational machines that can learn on their own, generally with supervision, for this offers potential to obviate the long and laborious knowledge acquisition process. Machine learning technologies may enable unmanned vehicles to learn from their surroundings and mistakes. This would enable us not only to skip the initial knowledge acquisition process required to set up an AS, but also to reduce the amount of required initial knowledge. In other words, when a system can learn on its own, the burden of providing its initial knowledge can decrease dramatically, and the ability of the system to adjust to its unique mission-environment context and adapt to change can accelerate exponentially and proceed without human interaction.

As noted recently by Selby (2015), some robots can learn to recognize objects and patterns fairly well, but to interpret and be able to act on visual input is much more difficult. Nonetheless, researchers at the University of Maryland, funded by DARPA’s Mathematics of Sensing, Exploitation and Execution (MSEE) program, recently developed a system that enabled robots to process visual data from a series of “how to” cooking videos on YouTube. Based on what is shown on a video, robots are able to recognize, grab and manipulate the correct kitchen utensil or object and perform the demonstrated task with a high degree of accuracy—without human input or programming.

“The MSEE program initially focused on sensing, which involves perception and understanding of what’s happening in a visual scene, not simply recognizing and identifying objects,” said Reza Ghanadan, program manager in DARPA’s Defense Sciences Offices (Selby, 2015). “We’ve now taken the next step to execution, where a

robot processes visual cues through a manipulation, action-grammar module and translates the cues into actions.”

Another significant advance to come out of the DARPA research is the robots’ ability to accumulate and share knowledge with others. Current sensor systems typically view the world anew in each moment, without the ability to apply prior knowledge. “This system allows robots to continuously build on previous learning—such as types of objects and grasps associated with them—which could have a huge impact on teaching and training,” Ghanadan said (Selby, 2015). “Instead of the long and expensive process of programming code to teach robots to do tasks, this research opens the potential for robots to learn much faster, at much lower cost and, to the extent they are authorized to do so, share that knowledge with other robots. This learning-based approach is a significant step towards developing technologies that could have benefits in areas such as military repair and logistics.”

In his article on MSEE, Dr. Reza Ghanadan (Ghanadan, 2011) also notes the following.

The goal of the Mathematics of Sensing, Exploitation and Execution (MSEE) program is to explore and develop high-impact methods for scalable autonomous systems capable of understanding scenes and events for learning, planning, and execution of complex tasks. The program is exploring powerful mathematical frameworks for unified knowledge representation for shared perception, learning, reasoning, and action. One of the central concepts in MSEE is to exploit methods based on minimalist generative grammar, similar to human language, to represent and process visual scenes and actions. Data-driven methods for spatial, temporal, and causal parsing of information are being developed for semantic understanding of scenes and events in unstructured environments along with cognitive processing methods for exploitation and manipulations. The foundational premise of the program is that a comprehensive mathematical framework to describe an integrated SEE system would allow for detailed analysis of its potential performance and serve as a guide to prototype design. Methods will be demonstrated in use cases motivated by defense applications

such as intelligence, surveillance and reconnaissance (ISR) and vision-guided robots to perform repairs.

The MSEE program aims to address a growing difficulty: The amount of data collected by DoD sensor systems exceeds the ability of human analysts and current automated decision systems to extract actionable information. This data deluge is pervasive throughout the military and applies to single and multi-modal sensing platforms. Today, evaluation methods rely on feature detection and category classification using individual pipelines for different tasks due to the lack of an effective unified representation. Hence, under the current paradigm, the semantics derived from sensor outputs do not emerge until an analyst has assimilated the data.

As a result of this dynamic, three challenges emerge:

data's worth can only be evaluated after analysts have interpreted it, and with knowledge of how it will be used;

prior knowledge, including that which may have accrued during previous processing, is not used during the production of sensor products; in other words, sensors process signals as if they were seeing the world anew at every instant; and,

in the presence of multiple sensors, analysts must reconstruct a unified view out of sensor output products that were not, in general, designed for integration.

A new approach to sensing is required to confront these shortcomings. An effective representation for recognizing objects, attributes and actions, and for parsing spatial-temporal relational information, would result in scalable platforms capable of autonomous learning, inference, reasoning, planning and execution of complex tasks.

The goal of the Mathematics of Sensing, Exploitation and Execution (MSEE) program is to capture the economy and efficiency that would derive from an intrinsic, objective-driven unification of sensing and exploitation. The foundational premise of the program is that a comprehensive mathematical framework to describe an integrated SEE system exists. Such a theoretical

description would allow for detailed analysis of its potential performance, serve as an invaluable guide when constructing a prototype to demonstrate the effectiveness of the approach, and enable quantitative determination of the effective utility of various sensors and sensing modalities.

MSEE includes three planned phases. The goal of Phase I is to create the mathematical foundation for a representation-centric model. The goal of Phase II is to refine the representation constructed in an initial software system able to answer queries related to the content of sensor data from a single modality. The goal of Phase III is to develop a fully integrated, modular system that demonstrates quantitative and qualitative analyses of systems that integrate sensing/perception, exploitation, and execution, with multi-modal sensor data. ... To date, performers on MSEE have pursued fundamental research into the nature of stochastic modeling and knowledge representation. These basic tools are being used to build prototype systems. Progress in these areas should greatly advance DoD's ability to build high-performance systems in a number of areas including ISR and supervised robust autonomous systems.

3. Biologic Design

U.S. intelligence experts are trying to reverse engineer the algorithms of the human brain by blending data science and neuro science in attempts to make rapid advances in machine learning and artificial intelligence (Keller, 2015). To the degree that natural language analytics can play a role, Watson¹⁰ might be useful. The key elements and functions of this advanced technology include natural language processing, hypothesis generation and evaluation, and dynamic learning. Sometimes it's used in conjunction with other analytics (such as real-time data analytics) to inform decisions (Codella, 2012).

Current initiatives are focused on creating ever more capable, neurally derived machine learning algorithms, with an aim to improve the ability to perform complex information processing tasks such as one-shot learning, unsupervised clustering, and scene parsing with near-human proficiency.

¹⁰ Watson is a cognitive system designed and built by IBM.

IARPA (Intelligence Advanced Research Projects Activity) computer and information technology researchers are intrigued with the data processing abilities of the human brain, specifically its ability to separate and categorize signals robustly in the presence of significant noise and non-linear transformations, and then extrapolate and apply single examples to entire classes of stimuli (Keller, 2015).

Air Force researchers seek to automate the collection and use of intelligence information gathered from many different platforms and correlated in several different ways, make better use of raw sensor data from existing multisource, multiplatform, real-time collection systems, and automate intelligence information processing for assessment, cueing, electronic attack, and battle damage assessment (Keller, 2010).

Many say that today's state of the art algorithms are brittle and do not generalize well—at least nowhere near the level of the human brain. Today's leading machine learning algorithms and the human brain still differ significantly in the details of their operation.

Further, Johnson et al. (2004) describe the design, development and testing of a UAS with highly automated search capabilities. These capabilities allow the system to search a prescribed area, identify a specific building within that search area based on a small identifying sign located on one wall, and locate a potential opening into that specified building. Subsequent to selection of the search area, all functions are automated and do not require human operator assistance. The applications of these capabilities include reduction of operator workload and facilitation of guided-munition missions, conducted without the assistance or intervention of human operators.

This work builds upon previous development of a research UAS and image processing algorithms, and it is the first publication of the method used for automated mission management (i.e., the automation of mission-level decisions). Of particular significance is the fact that this work was carried to the flight test phase and was tested under realistic conditions. This introduces all of the issues relating to real-time algorithms, dealing with noise/clutter/uncertainty, and logistics which are all so important to practical automated mission management.

Researchers anticipate achieving a quantum leap by creating machine learning algorithms that use neurally inspired architectures and mathematical abstractions of the

representations, transformations, and learning rules employed by the brain. Targeting neuroscience experiments that interrogate the operation of mesoscale cortical computing circuits, and taking advantage of emerging tools for high-resolution structural and functional brain mapping, offer excellent promise, which should facilitate iterative refinement of algorithms based on a combination of practical, theoretical and experimental outcomes.

4. Computational Hardware

Big firms like Amazon, DHL and Google are developing their own drone fleets for rapid delivery of consumer goods, fast food and pharmaceuticals. However, current FAA rules restrict drones to flying within visual range of human operators because of the risk of collision. Drones need an automatic “sense-and-avoid” capability before they will be allowed to make deliveries on their own.

Computers and their integrated sensor suites must be capable of recognizing objects in a timely manner in order to provide sufficient warning to avoid collisions. The key may be to mimic how brains work; our brains are poor at number-crunching but can process complex sensory input faster than digital systems.

As noted by David Hambring in his New Scientist magazine article (Hambring, 2015), Bio Inspired Technologies of Boise, Idaho, is building a sense-and-avoid system using a memristor, which integrates a resistor with a memory into a single device. Like the synapse in a biological brain, the memristor changes when impulses pass through it. Crucially, it is able to remember the impulse after it has stopped. This capability forms the basis of a learning system that mimics neurons and the connections between them. A chip-sized neural system linked to the drone’s existing camera can be trained to recognize aircraft and other hazards at long range. Bio Inspired’s drone should be ready for its first flight in the very near future.

The system can also recognize objects like clouds, birds, buildings and radio towers, and it uses visual cues to estimate how far away the objects are. “Objects like other aircraft can be catalogued in a vague sense, meaning ‘I see an aircraft’, or in an exact sense: ‘I see another drone’,” says Terry Gafron, CEO of Bio Inspired (Hambring,

2015). Equipped with this information, the drone plots a new flight path to avoid a potential hazard and updates the hazard's position in real time as the threat moves.

"Nature seems to use this approach very effectively," says David Warne of Queensland University of Technology in Australia, who has worked with artificial neural networks that allow drones to recognize vegetation (Warne, 2014, page 5). Like others in this area, much of Bio Inspired's research has been funded by the military, but it offers the potential to serve a much broader market. A robust sense-and-avoid capability will make it possible for fleets of small drones to crisscross cities delivering packages, for instance. Like a bird or insect, a neural-enabled drone could fly to the trickiest landing places, even balconies.

In the industrial field, neural drones could patrol pipelines looking for leaks, or identify electrical faults on power lines. Closer to home, smart drones could clean windows, pick up litter, clear gutters, weed your garden, or send information to your car about which parking spaces are open. "It simply flies around town monitoring parking spaces," says Gafron (Brown, 2015). Smart drones could even track animal populations, flying along livestock boundaries to track wolf populations, for example. "Not only could the system fly autonomously, but it could conceivably tell the difference between a deer and a wolf from the air," Gafron says.

Memristor-inspired drones are not the only approach. In 2014, a US organization, DARPA, (Defense Advanced Research Projects Agency) unveiled the TrueNorth neural chip developed in conjunction with IBM. This project is a simulation of a neural network using digital hardware with enough neurons to match agile flyers like bees. This article appeared in print under the headline "A drone that learns"

5. Perception

The popular term "drone," which conjures images of remote-controlled flying zombies, is becoming less and less descriptive of the latest UAS technology. New applications are requiring more autonomy and intelligence. "When people think about drones, they largely think of big military assets that are flying high in the sky where there's not a whole lot of anything to hit," says Nick Roy, director of the Robust Robotics Group at MIT's Computer Science and Artificial Intelligence Laboratory (Brown, 2015). "But there are a lot of applications for smaller scale UAS working closer to the ground

that require more autonomy, such as agricultural monitoring, package delivery, and situational awareness for first responders."

Teaching UAVs and other robots to think for themselves is the central mission of the Robust Robotics Group. "We want UAVs to be able to operate in urban environments, to get useful things done, and interact with people," says Roy, who is also an associate professor of aeronautics and astronautics (Brown, 2015). "We want them to become as intelligent as they need to be for the task at hand."

Roy has recently focused more on UAVs than terrestrial robots, although many of the principles and algorithms are similar. UAVs will require more autonomy to avoid collisions and crashes, as well as to understand what's happening around them. Some level of reliable autonomous operation will be essential if the FAA is to fully permit commercial applications in the United States. "It's not just about avoiding obstacles, but about understanding the environment and what's safe and unsafe," Roy says. "UAVs need to understand their own behavior in terms of reliability and performance, and also to understand how people want them to do things" (Brown, 2015).

Indeed, Roy accepted a sabbatical position in 2012 at Google X to help launch Project Wing, a project with the goal of demonstrating the viability of product delivery using UAVs. Project Wing is a hybrid aircraft instead of the typical quadrotor designs that have recently dominated the academic research and the consumer UAV market. Although it does use four rotors, the rotors normally perform like airplane propellers. When the craft reaches its target to drop a package, it tilts upward so it can hover like a quadrotor. This "tail sitter" design is a revision of an old idea that has yet to be proven commercially feasible. "Hybrid vehicles like tail sitters, tilt rotors, tilt props, or vehicles with two propulsion systems, have been explored throughout the history of aviation," Roy says. "But enough things have changed to make them worth trying again. Our ability to manufacture small vehicles and put computation and modern control systems onboard means the things that once were hard are relatively easy now" (Brown, 2015).

Compared to quadrotors, conventional fixed wing craft have obvious limitations, including the need for a runway and the requirement for minimum airspeed to remain airborne, Roy says. Yet, "fixed wing craft are a lot more efficient in flight and can stay up much, much longer," he adds (Brown, 2015).

Although Roy is focused more on software than hardware, he must keep abreast with all of the latest technologies, especially sensors, which help shape the way the UAV thinks. Spurred on by the need to reduce weight and power consumption, for example, some UAV researchers are aiming to use lightweight, low-cost cameras for navigation, rather than requiring LIDAR equipment or 3-D cameras. "Passive cameras give you an understanding of the scene that I think will be important in the future," Roy says. "Pure vision-based navigation has yet to work reliably, but the field has progressed a lot. I'm excited about how we might use passive cameras to help UAVs navigate on their own" (Brown, 2015).

In the meantime, no single sensor technology is the right answer, Roy says. "GPS has issues in urban environments and cameras have issues especially at night," he says. "A lot of my group's recent research has focused on accurate ranging, whether it's a laser range finder or a 3-D camera. Those sensors are heavy and don't work in every domain. The right answer will probably lie in a fusion of sensors" (Brown, 2015).

Delivery UAVs like Project Wing or Amazon's prototype will need more autonomy and intelligence than a typical UAV used for crop monitoring or filming commercials. This is especially true if the UAV is expected to drop off and pick up packages in urban environments. "The UAV will need to be smart enough to reason about its own performance and impending failures," Roy says. "Autonomy is the biggest challenge facing integration in the airspace. Vehicles need autonomy in order to recover from failures, and to see other aircraft and not hit them. They need autonomy to interact with air traffic control and play nicely in the national airspace" (Brown, 2015).

Researchers at MIT and elsewhere have focused on imbuing robots with object recognition, but that's only the beginning. A greater challenge is to bridge the gap between the fundamentally different ways in which people and robots think. "Robots think about the world in terms of very low level geometry," Roy says. "They don't think of walls as walls, but rather as pixels they can't drive through. To work with people, robots must understand what things are for. To ask a robot to collect a box or load a truck, it needs the semantic understanding of what these objects are" (Brown, 2015).

Roy is focused less on object recognition itself than on helping robots "understand how objects are distributed and how they can interact with them," he says. "Once you

have object detection or scene understanding, you can move to the next step: showing the robot how to use this understanding to make decisions" (Brown, 2015).

One of Roy's students, for example, is attempting to improve a UAV's understanding of wind patterns in the urban environment. The UAV could then use that knowledge to avoid turbulence or choose minimal energy routes. Despite the continuing advance of computer miniaturization, the weight and power limitations of UAVs will continue to challenge their ability to process information quickly enough to make timely decisions. Rapidly fusing and integrating data from multiple sensors poses "computational challenges that are outside the scope of real-time systems like UAVs," Roy says. "A lot of my research involves finding useful approximations that involve getting very good answers at the cost of a little accuracy and precision" (Brown, 2015).

These approximation algorithms were put to work in Roy's recent experiments in which a fixed wing vehicle carrying a laser range finder flew at nearly full speed around the tightly constrained environment of a parking garage. "If you were to try to incorporate the laser range finder into the full-state estimate of the vehicle's 12 degrees of freedom, the computation would get intractable," Roy says. "But if you break the problem apart into the bits that the laser finder can see at any one time, you can still get the right answer, but much more efficiently than if you ask the laser to 'reason develop' the entire system at once" (Brown, 2015).

Teaching UAVs to recognize objects and process sensor data in order to make real-time decisions will help avoid collisions even in complex environments, including offices. Yet, additional autonomy and intelligence is required when UAVs work closely together with people. Beyond ensuring the safety of humans, the algorithms need to be sufficiently sophisticated to enable UAVs to take instructions from people or collaborate with them to get things done.

C. PROGRAMMATICS

In this section we elaborate our analyses and insights pertaining to programmatics. Despite compelling reasons to employ more autonomous functionality in AS, a variety of inter-related reasons contribute to a collective inability in terms of providing such capabilities.

For one, by their very nature, autonomous capabilities are enabled by software. The proportion of software costs to overall system costs continue to grow with the increasing complexity of the software itself. Additionally, a brief examination of the literature reveals numerous autonomous functions have been developed by many groups; yet little of this capability finds its way into operationally fielded systems. Examining commercial applications point to several approaches that may help address this issue.

First and foremost, the DoD must recognize that it is no longer the biggest player in the unmanned systems field. Soon that organization will assume a similar role to where it is in computers. DoD may have initiated this technology, but now it is having a smaller impact. In a world where you can order a drone on Amazon for \$200, and the largest manufacturer of drones is the China based DJI, the DoD has no choice but to adapt.

When trying to determine when these autonomous capabilities will become pervasive, it is insightful to consider the automobile industry. Many people are awaiting the arrival of the driverless car (like the one Google is driving on California roads) and wondering when they will be able get autonomy in their cars. Meanwhile, several of the autonomous capabilities that will be standard in most driverless cars already exist in their cars today. These include ubiquitous technologies like GPS and ABS and recent developments such as lane change detection; driver alertness monitoring; front crash prevention systems; and automatic parallel parking systems.

Autonomy seems unlikely to be developed in one big bang. Rather, it will more likely evolve from basic autonomous functional components. Although such components must be integrated into a system, they will be supplied from a variety of sources. For many types of components, the manufacturers are relatively stable and have well established ties to DoD. In the area of software that provides autonomous functionality, however, relatively few coincide with classic defense contractors, and the DoD may have to adjust to new suppliers more than vice versa.

The typical model of defense acquisition appears ripe for change to accommodate autonomous systems, and open architecture (both technical and programmatic) appears to be key in this regard. For instance, open architecture enables users to rapidly integrate autonomous capability components on systems throughout their life cycle. The DoD Open Systems Architecture Contract Guidebook (DoD Data Rights Team, 2013) provides

insight and guidance, but programs with fielded systems will need to improve in terms of accountability, to evolve to modular open system architectures for autonomy, and to develop acquisition strategies that allow easier third party integration. This will likely require such programs to take more of a lead system integrator role.

An example of a program attempting a more open approach is the Common Control System (CCS) for Naval UAS (NAVAIR News, 2011). Following product integration, which will ultimately optimize UAS interoperability and sensor functionality, the CCS will provide flexibility and allow efficient services and control, which is currently an immature technology. Another significant effort, initiated by NAVAIR, is the Future Airborne Capability Environment (FACE; see Adams, 2014). The FACE initiative, which is a combined US Army and Navy program, will ultimately standardize avionics software for UAS, slash acquisition costs, and reduce time to market.

Another programmatic challenge is providing capability across the metaphoric “valley of death” between science and technology (S&T) research efforts and fielded systems. In addition to providing open architecture environments for the S&T community, overcoming this challenge will require programs to be held accountable for pulling in promising autonomous S&T capabilities as they emerge. For instance, a developer focused on autonomy could be required to conform to a modular design that aligns with current open architectures of fielded systems or systems under development.

Integrating autonomous functional components from a variety of sources that are in continuous competition is just the first step. Typically, testing and certifying software is the most costly and time consuming aspect of the development. In the past, developers have attempted to exercise all possible logic branches during the test/certification phase, particularly for safety and/or security critical functions. However, as autonomous behaviors continue to mature, this approach is becoming unaffordable.

Moreover, as AS developers continue integrating ever more non-deterministic functions (e.g., autonomous algorithms that learn and change behavior based on what they learn), it will become infeasible to test all possible outcomes. Instead, new designs, system analyses, cost-efficient testing and certification processes will all become necessary, as will tools to allow certification of such algorithms while maintaining safe

and secure operations. Designs will need to reflect architectures that facilitate appropriate segregation and robustness to be reliable, safe and secure.

Indeed, algorithm and software certification processes may begin to look more like those used to certify people. As with people, they are given more responsibility with experience only as they demonstrate that they can perform important tasks in a reliable and appropriate manner. As with people also, for another instance, since they learn and change continually, algorithm and software certifications will need to be updated on an ongoing basis. This will necessitate tools for modeling and simulation, with fidelity spanning the range from basic to high, in addition to flight test platforms with open architectures that allow the routinely use of surrogate vehicles prior to integrated testing on an operational platform.

D. EDUCATION AND TRAINING

In this section we elaborate our analyses and insights pertaining to education and training. It is important to realize that the area of autonomy is highly dynamic. Change is fast and very difficult to anticipate. This highlights the importance of education and training in Science, Technology, Engineering and Math (STEM), particularly as it focuses on robotics and autonomy. Competitions such as FIRST, VEX, and those sponsored by the AUVSI Foundation give young STEM students a level of excitement on par with sports teams. FIRST (For Inspiration and Recognition of Science and Technology) is an international youth organization that operates robotic competitions and fosters the development of STEM skill sets in today's younger generations. VEX is an organization that sponsors world-wide robotics competitions primarily for middle school and high school students. The AUVSI (Association of Unmanned Vehicle Systems International) Foundation, with significant financial support from the Office of Naval Research, sponsors UAV, USV and UUV (ROBOSUB) competitions primarily for US and international high school and college students.

DoD labs could go further, for instance, by partnering with regional academic institutions that are working in these areas, and seek to enhance learning. STEM graduates are highly valued in the commercial world, and the DoD will likely face stiff

competition for such talent. A focus on education and training may indeed be as important as our current emphasis on technology.

E. CULTURE

In this section we elaborate our analyses and insights pertaining to culture. The cultural acceptance of autonomy is clearly one of the most challenging issues. We have identified concerns with trust and ethical issues associated with the employment of unmanned systems. Will people learn to trust machines to think and act for them? Can machines be taught to make ethical decisions? Such questions point away from technology—although they have technologic underpinnings—yet are as or more important than the technology itself. What benefit can one expect to derive from integrating manned and unmanned aircraft (e.g., in common squadrons, on common missions, through training exercises) if human pilots refuse to trust or even fly with their machine counterparts? The greatest autonomy technology in the world will fail to support TASP if people cannot accept and learn to leverage this robust capability.

"We need to teach robots how to interact with people as seamlessly as people work with each other," Roy says. "They need things like semantic maps to help them think about the world the same way people do. They also need to understand what people want and how they behave. We're looking at things like natural language interfaces, and connecting human speech with the things the robot sees" (Doshi, 2007).

Toward this end, the Robust Robotics Group has made some progress in teaching robots to understand directions and instructions. Now, the group is working on dialog management: teaching the robot and human how to converse. "The challenge for the robot is not just how to know it needs to ask a question, but how to ask the question in a way that can return a useful answer," Dr. Nicholas Roy says. "If the robot says, 'I don't understand,' the human will probably get annoyed and abandon the robot. ... This technology has to mature substantially before we're really ready to have robots become part of our everyday lives" (Doshi, 2007).

We also find considerable fear and criticism regarding lethal autonomy (e.g., arming AS). Heather Roff, for instance, asks how too much autonomy could possibly impact public perception and once again offer an ethical dilemma (Roff, 2015), while

others fear that autonomous weapons could spark an arms race that would increase both the likelihood of wars and the slope of their escalation (Roff, 2015). We have also identified an open letter from the Future of Life Institute (with 16,000 signatories, including Stephen Hawking, Steve Wozniak and Elon Musk), which calls for a ban on artificially intelligent weapons. Civilian politicians, who are elected by the voting public, sit atop every military organization in the US, so popular opinion matters.

F. OPERATIONS

In this section we elaborate our analyses and insights pertaining to operations. Autonomy is important for operations in many ways. One is to address the key vulnerabilities of AS. For instance, autonomous systems need to operate in a communication (including GPS) limited or denied environment. Bandwidth is limited even under the best of circumstances, and loss of all communications for a period of time is frequent. To cope with this limitation at present, a UAS will typically exhibit a return to base behavior, or it will initiate a lost communications mode, both of which are preprogrammed and enabled automatically.

However, future systems will need to be more capable. One can't rely on 100% communications when operating routinely in the National Air Space (NAS), for instance. According to the CNO SSG "If the Navy could overcome the communication and precision navigation limitations, there could be nearly a four-fold increase in the situation where unmanned systems could carry out the missions" (Hogg, 2009). Indeed, AS will be required to operate in a "network optional" manner. When a patrol of warfighters is sent on a mission, for instance, they will be most effective when they have communications that permit them to exploit intelligence from other sources; coordinate their actions with other units; and receive commands when required. However, they must also be capable of maintaining effectiveness even if this communication capability is lost or degraded.

Another operational reason for autonomy centers on functions where computers can outperform people. For instance, it would be extremely difficult for a human pilot to remotely control an air vehicle and perform a consistent safe landing in the highly dynamic shipboard environment. Hence the US Navy has successfully demonstrated that

it can employ autonomous shipboard recovery algorithms for a variety of UASs including the MQ-8B, MQ-8C, X-47B, RQ-21A, and the Scan Eagle.

Perhaps the most pressing operational need for autonomous functions, in our AS, is to address a problem that is associated with one of their primary strengths; that of persistence. At a mission radius of 2000 nm, the Triton (RQ-4) UAS (three vehicles), for instance, is designed to provide persistent ISR, 24 hours a day, 7 days a week, at least 80% of the time. The Triton employs a variety of sensors including a 360-degree Field of Regard (FOR) Multifunction Active Sensor (MFAS) electronically steered array radar, an Electro-Optical / Infrared (EO/IR) sensor, an Automatic Identification System (AIS) receiver, and Electronic Support Measures (ESM). The amount of raw data produced by this system is staggering. The bandwidth required to transmit this volume data is huge. Sophisticated autonomous algorithms, which would convert data into actionable intelligence onboard the aircraft, would simultaneously improve the effectiveness of the system; decrease its communication vulnerability, and reduce the number of humans required to deal with this large amount of data. Naval forces, whether sea based or a Marine Expeditionary Unit, do not have the capacity to add large numbers of personnel to operate and exploit AS.

Another area where autonomous capabilities will prove invaluable is through evolution from the current paradigm of one or more “operators” per vehicle to a system of systems approach, in which a small number of people oversee multiple vehicles and sensors as mission managers, and in which the vehicles manage themselves. This would decrease the manning required for UAS operations, and with sufficient progress in terms of autonomy enable swarming behaviors, in which a large number of relatively inexpensive systems could autonomously collaborate to overwhelm an adversary. This concept was recently demonstrated by the ARSENL (Advanced Robotic Systems Engineering Laboratory) team at NPS, for instance, where 2 operators managed 50 small UAS.

This mission manager concept would also permit exploitation of multi-domain AS in a manner that individual systems can’t effectively achieve. As a current example, consider the Explosive Ordnance Disposal (EOD) mission, for which the Navy is the lead service. Currently an EOD team must be deployed to the location of a suspected

Improvised Explosive Device (IED) or like hazard. The team is exposed to possible enemy fire while transiting to and from the IED site and while on site. The team typically employs an Unmanned Ground Vehicle (UGV) in line of sight teleoperation mode.

However, as a use case, consider if the system were to include a UAS (e.g. Fire Scout) in conjunction with the EOD UGV. The IED mission could potentially be managed without any EOD personnel at the suspected IED site. For instance, the UAV could carry the UGV to an area of interest, survey the area to determine the optimal placement of the UGV, and provide ideal tactical path information from the UGV to the IED location. It could also act as the communications relay between the UGV and the mission commander as well as provide overwatch surveillance of the mission area looking for possible ambushes and awareness of civilians potentially at risk. The mission could be managed seamlessly from a single control station, and tasks related to the vehicles could be managed by onboard autonomous software. The UAV could then retrieve the UGV upon rendering the IED site safe and ferry it back to base. The mission as such could be accomplished faster, more effectively, and without risk to personnel.

Another operational use case centers on whether an autonomous UAS (e.g., unmanned helicopter) might be more effective than people at making targeting decisions. In the heat of battle, for instance, a soldier may be tempted to return fire indiscriminately, in part to save his or her own life. By contrast, a machine (today) won't grow impatient or scared, be swayed by prejudice or hate, willfully ignore orders, or be motivated by an instinct for self-preservation. Indeed, many researchers argue for speedy deployment of self-driving cars on similar grounds. Vigilant electronics may save lives currently lost because of poor split-second decisions made by people. How many soldiers in the field might die waiting for the person exercising "meaningful human control" to approve an action that a computer could initiate instantly?

G. STRATEGY

In this section we elaborate our analyses and insights pertaining to strategy. Today, Pentagon officials are intensifying their quest for technologic superiority. It "is one of the most important strategic tasks and risks facing our Department," Deputy Defense Secretary Bob Work said at the opening of the China Aerospace Studies Institute

(Work, 2015). “Because if we allow our technical superiority to erode too much, again, this will undermine our conventional deterrence. It will greatly raise the cost, the potential cost of any intervention overseas, and will contribute to crisis instability.”

Deputy Work expands upon his vision of future autonomy in war. “You’ll have a high degree of human-machine collaboration, like free-style chess, in which machines, using big data analytics and advanced computing, will inform human decision makers on the battlefield to make better decisions than humans can do alone or machines can do alone,” he said (Work, 2015). “You’re going to have routine manned and unmanned teaming. You’re going to have increasingly capable autonomous unmanned systems. You are going to have all of this. So the future of combat, we believe is going to be characterized by a very high degree of human-machine symbiosis, such as crude platforms controlling swarms of unmanned, inexpensive unmanned systems that can be flexibly combined and fielded in greater numbers.”

These types of systems, far beyond the cheap and immature ones of today, herald a qualitative shift in strategy. If the U.S. begins to wage war in this way—or even gains the ability to do so—then competitor nations will likely grow fearful and feel vulnerable. This is a strategic concern in two aspects regarding the kind of damage that can result from an arms race between nations. First, arms races—as interactive competitions between rival states, where the competitors build up particular weapons technologies, capabilities or personnel over time—increase not only the probability of militarized disputes between competitors, but also the probability of escalation when those disputes erupt. Arms races make war more likely and more violent.

Second, the type of technology we are discussing here extends well beyond conventional weapons in conventional war. We are talking about creating weapons that push the boundaries of artificial intelligence and autonomy. This push to create adaptive, learning and intelligent weapons platforms will ultimately require greater onboard capabilities, less communication, and a system-of-systems approach to war. They will also effect delegated decision-making and accelerate the tempo of war. Quite central to the thrust of our research stream, this will challenge—and very quickly break—the current C2 structure. The greatest technologic accomplishments in terms of autonomous

systems are useless unless they can be commanded to do what is desired and controlled to avoid doing what is not desired.

Moreover, the metaphoric price of entry into AS is relatively low. Most unlike the nuclear stockpiling that occurred during the Cold War, and quite different from large-scale major weapon systems of today (think aircraft carriers, jets and tanks), autonomous systems are comparatively very inexpensive to acquire and straightforward to develop and deploy. One can purchase a credit-card sized Raspberry Pi computer for \$30, for instance, and a \$200 quad copter can be bought online. In this way, autonomy may accentuate asymmetric warfare, giving smaller, even non-state adversaries access to disruptive technologies. This is in parallel to the manner in which cyber attacks, IEDs and like, inexpensive capabilities are affecting the tactics and disrupting the operations of even the richest nations and most powerful military organizations in the world.

THIS PAGE INTENTIONALLY LEFT BLANK

V. CONCLUSION

The technologic capabilities of autonomous systems (AS) continue to accelerate, and integrated performance by AS and people working together can be superior to that of either AS or people working alone. We refer to this increasingly important phenomenon as Teams of Autonomous Systems and People (TASP), and through our recent research—representing the current state of the art—we have demonstrated a computational experimentation capability in the TASP domain. The problem is, several technology trends suggest that unmanned aircraft may be diverging away from operating and behaving like their manned counterparts, suggesting that some of our most futuristic model assumptions may be off target.

This is where our ongoing research project continues to make an important contribution. The project described in this technical report centers on expanding our recently enabled C2 computational modeling and simulation capability to understand the next generation of unmanned aircraft systems, with particular emphasis on specifying advanced models for computational experimentation and the potential influence of Artificial Intelligence. In particular, we investigate technology trajectories and design visions for next generation unmanned aircraft systems through qualitative methods.

In this technical report, we provide an overview of the POWer computational experimentation environment and summarize results from our experiments on TASP C2 organizations and phenomena, for these constitute key background areas of interest for the current study. Such results fall in three primary areas: 1) Autonomy Degree, 2) Interdependence Levels, and 3) C2 Implications.

Then we summarize the qualitative research method. Specific techniques include archival review, semi-structured interviews and participant observations. Qualitative data are collected, coded and analyzed continuously, with the data collection effort not stopping until theoretical saturation is reached. Coding adheres to well-accepted grounded theory building techniques, including open, axial and selective coding activities in series.

Key results follow, which we organize into the six primary themes that emerge from our coding and analysis of qualitative data: 1) technology, 2) programmatic, 3)

education and training, 4) culture, 5) operations, and 6) strategy. Within the technology theme, we identify five major, complementary, interrelated drivers of increasing autonomy. These include artificial intelligence, machine learning, biologic design, computational hardware, and perception. We cite a number of examples and challenges for each driver, and we explain how they affect autonomy today as well as projecting how they portend to continue doing so in the future.

One major implication is that such complementary technologies combine to enable degrees of autonomy that far surpass systems in operation today. Another is that autonomous systems are becoming increasingly capable in their behavior, yet we find a diversity of technologic approaches to enabling such behavior. Traditional AI with its long and laborious knowledge acquisition process is giving way to machine learning techniques, and researchers are leveraging the human brain and biologic systems for technical insight, inspiration and architecture pertaining to autonomous thinking and learning. A third implication centers on computational hardware such as the memristor, which has biologic inspiration and enables new integrated capabilities, and the ever expanding requirement for UAS in particular to perceive their environments.

Within the programmatic theme, we note several reasons why autonomous capabilities remain slow in terms of reaching the field for operations. Such reasons center on the software intensive nature of such capabilities and the challenges associated with software development and acquisition. We also note how numerous software developers in this area are not among the large DoD focused firms that build airplanes, ships and tanks; rather, many are comparatively very small firms, for which the DoD does not represent a major customer. Open systems architectures are very important too, and we note the challenges inherent in attempting to certify advanced autonomous systems, which cannot be tested exhaustively, particularly as they continue to learn and adapt after operational release and use.

The education and training theme emerges as a very important one, in which we highlight the importance of science, technology, engineering and math (STEM), particularly as it focuses on robotics and autonomy. Indeed, a focus on education and training may be as or more important than our current emphasis on technology.

Likewise with the culture theme, for the cultural acceptance of autonomy is clearly one of the most challenging issues. We have identified concerns with trust and ethical issues associated with the employment of unmanned systems. Such concerns point away from technology—although they have technologic underpinnings—yet are as or more important than the technology itself. What benefit can one expect from integrating manned and unmanned aircraft (e.g., in common squadrons, on common missions, through training exercises), for instance, if human pilots refuse to trust or even fly with their machine counterparts? The greatest autonomy technology in the world will fail to support TASP if people cannot accept and learn to leverage it. Hence a very important research thrust emphasizes teaching autonomous systems how to interact with people as seamlessly as people work with each other. We also note how public fear of armed autonomous systems has the potential to undermine even heroic technologic progress through the democratic voting system and civilian leadership of military organizations in the US.

Within the operations theme, we identify numerous operational needs for autonomy. Addressing current vulnerabilities in terms of interrupted communication represents one notable example, but leveraging the capabilities of computers that dominate human performance represents a critical example that undergirds TASP. Reducing the manning levels required to operate UAS is clearly important too, where the mission manager approach could enable AS swarming and like behaviors.

Finally, the strategy theme provides an opportunity for us to summarize aspects of autonomy which are causing strategic disruption and need for change. Military leaders at the highest levels recognize the potential of AS in warfare and for deterrence alike, and the need for people and machines to work cooperatively together rises to the top of critical needs in this regard. We also note the potential for new arms races based on AS capabilities, and we emphasize the asymmetric nature of AS: the metaphoric price of entry into AS is relatively low.

Here we conclude in turn with our agenda for continued research along these lines. Through the present investigation, we gain greater insight into and understanding of autonomy trends—not just technologic trends, but programmatic, educational, cultural,

operational and strategic—that inform our vision of TASP over the coming five to ten years.

Instead of pragmatically assuming that the future will be as we have been envisioning it (esp. that unmanned systems will evolve to operate and behave increasingly like their manned counterparts), now we have several, rich, interrelated themes to help organize a definitively multifaceted future, and we can see some important aspects of TASP that we had not deemed previously to warrant inclusion in our models (e.g., programmatics, education and training, culture, strategy). We also now have a finer grained view of technological advancements, which we can grasp more easily in terms of complementary, interrelated drivers such as AI, machine learning, biologic design, computational hardware and perception.

The next step is to understand the implications of this new knowledge in terms of computational modeling and to refine our POWer models to support correspondingly finer grained and more knowledgeable computational experiments on TASP C2. Integrating and building upon results from our previous computational experiments will provide an excellent basis for progress in this regard, especially as we have a number of novel C2 organizations and approaches (e.g., Collaborative, Self-Synchronization) to assess. This represents the research trajectory on which this present study falls, and we welcome other researchers, leaders and policy makers to join our effort to develop good answers and to provide effective guidance. C2 represents the single, most important determinant of military efficacy (Nissen, 2013), and although we have been learning and mastering C2 over many millennia, we're witnessing quantum change in terms of AS at present, and the military that masters the associated C2 most quickly may very well win the next encounter involving TASP.

VI. REFERENCES

Adams, C. March 2014. *FACE Software Effort Builds Momentum* Avionics magazine available at: aviationtoday.com/av/military/FACE-Software-Effort-Builds-Momentum_81308.html#.VggogM_luos.

Alberts, D.S. & Hayes, R.E. 2003. *Understanding: Command and Control* Washington, DC: Command and Control Research Program.

AMoD. 2001. Aviation history as Global Hawk completes US-Australia flight. Australian Ministry of Defence press release (24 April 2001); accessed from Wikipedia website: http://en.wikipedia.org/wiki/Northrop_Grumman_RQ-4_Global_Hawk ; retrieved on 27 January 2014.

Bamberger, R.J., Jr., Watson, D.P., Scheidt, D.H. & Moor, K.L. 2006. Flight Demonstrations of unmanned Aerial Vehicle Swarming Concepts, John Hopkins APL Technical Digest, 27(1): 41-55.

BBC. 11 July 2013. US drone lands on aircraft carrier. acc from Wikipedia website: http://en.wikipedia.org/wiki/Northrop_Grumman_X-47B#cite_note-BBC11July2013-4 ; retrieved 27 January 2014.

Bourne, D. May 2013. My Boss the Robot. *Scientific American* 38-41.

Brown, E. August 2015. *UAVs Learn to Fly Solo*. MIT News available at www.news.mit.edu.

Christiansen, T.R. 1993. Modeling Efficiency and Effectiveness of Coordination in Engineering Design Teams. Unpublished Ph.D Dissertation. Stanford University.

CFFC. 2014. *Fleet MQ-8B/C Fire Scout and MQ-4C Triton Unmanned Aircraft System (UAS) Platform Wholeness Concept of Operation*. US Fleet Forces Command.

Clear Path Robotics. 2009 TurtleBot 2 Open-source robot development kit for apps on wheels. Available at: <http://www.turtlebot.com/>

Codella, C. December 2012. *The Era of Cognitive Systems: An Inside Look at IBM Watson and How it Works*. REDP-4955-00.

Cohen, G.P. 1992. The Virtual Design Team: An Information Processing Model of Design Team Management. Unpublished Ph.D Dissertation. Stanford University.

Condon, B., Fahey, J., Rugaber, C., Lee, Y., Sterling, T. & Kurtenbach, E. January 24, 2013. *Practically human: Can smart machines do your job?* Associated Press.

Couts, A. September 12, 2012. Drones: 13 things you need to know from Congress's new report. *Digital Trends*; <http://www.digitaltrends.com/cool-tech/drones-congressional-research-service-report/#ixzz2fLx8ZLq4> ; retrieved 09/28/2012.

Defense Science Board – Task Force Report. 2012. The Role of Autonomy in DoD Systems. Available at <http://fas.org/irp/agency/dod/dsb/autonomy.pdf>.

DoDD 3000.09. November 21, 2012. Department of Defense Directive 3000.09, Autonomy in Weapon Systems.

DoD Data Rights Team. June 2013. DoD Open Systems Architecture Contract Guidebook for Program Managers v.1.1 available at https://acc.dau.mil/adl/en-US/631578/file/73333/OSAGuidebook%20v%201_1%20final.pdf.

Dolgin, D., Wasel, B. & Langelier, M. June 1999. Identification of the Cognitive, Psychomotor, and Psychosocial Skill Demands of Uninhabited Combat Aerial Vehicle (UCAV) Operators. Naval Air Systems Command technical report; DTIC code 19991004 162; retrieved 11/05/2012.

Doshi, F. & Roy, N. March 2007. Efficient Model Learning for Dialogue Management. Proceedings of the 2007 2nd ACM/IEEE International Conference on Human-Robot Interaction. Arlington, VA.

Feil-Seifer, D. & Mataric, M.J. 2005. Defining Socially Assistive Robotics. Proceedings of the 2005 IEEE 9th International Conference on Rehabilitation Robotics, Chicago, IL.

Galbraith, J.R. 1977. *Organizational Design*. Addison-Wesley.

Gateau, J.B., Leweling, T.A., Looney, J.P. and Nissen, M.E. June 2007. Hypothesis Testing of Edge Organizations: Modeling the C2 Organization Design Space. Proceedings International Command & Control Research & Technology Symposium, Newport, RI; Winner – Best Student Paper Award.

Ghanadan, R. February 2011. Mathematics of Sensing, Exploitation and Execution (MSEE), Available at <http://www.darpa.mil/program/mathematics-of-sensing-exploitation-and-execution>.

Glaser, B.G. & Strauss, A.L. 1967. *The Discovery of Grounded Theory: Strategies for Qualitative Research*, Chicago, Aldine Publishing Company.

Hambring, D. April 2015. *A drone that learns*. New Scientist Magazine issue 3017 Available at <https://www.newscientist.com/> and MIT News www.mit.edu/newsoffice/

Hogg. 2009. *The Unmanned Imperative*, CNO Strategic Studies Group XXVIII.

HRW. November 19, 2012. Human Rights Watch. Losing Humanity: The Case Against Killer Robots.; <http://www.hrw.org>; retrieved 11/26/2012.

Jin, Y. & Levitt, R.E. 1996. The Virtual Design Team: A Computational Model of Project Organizations. *Journal of Computational and Mathematical Organizational Theory* 2(3): 171-195.

Johnson, E.N, Proctor, A.A, Ha, J. and Tannenbaum, A.R.. December 2004. *Development and Test of Highly Autonomous Unmanned Aerial Vehicles*, Journal of Aerospace Computing, Information, and Communication Vol. 1.

Joyce, J.J. 5/10/2012. Scan Eagle UAS Deploying Aboard Coast Guard National Security Cutter, Navy.mil The Source for Navy News, Story Number: NNS120510-26; http://www.navy.mil/search/print.asp?story_id=67101; retrieved 01/19/2013.

Keller, J. March 2010, *Air Force to use artificial intelligence and other advanced data processing to hit the enemy where it hurts*. Military & Aerospace Electronics magazine.

Keller, J. January 2015, *Artificial intelligence and new levels of machine learning are aims of IARPA MICrONS program*. Military & Aerospace Electronics magazine.

Kunz, J.C., Levitt, R.E. & Jin Y. 1999. The Virtual Design Team: A computational model of project organizations. *Communications of the Association for Comp. Machinery* 41(11): 84-92.

Law, A.M. & Kelton, D. 1991. *Simulation Modeling and Analysis* 2nd Ed., New York, NY, McGraw-Hill.

Levitt, R.E. 2004. Computational Modeling of Organizations Comes of Age. *Journal of Computational and Mathematical Organization Theory* 10(2): 127-145.

Levitt, R.E., Orr, R.J. & Nissen, M.E. 2005. Validating the Virtual Design Team (VDT) Computational Modeling Environment. The Collaboratory for Research on Global Projects, Working Paper #25, 1-15. Available at: http://crgp.stanford.edu/publications/working_papers/WP25.pdf

Levitt, R.E., Thomsen, J., Christiansen, T.R., Kunz, J.C, Jin, Y. & Nass, C. 1999. Simulating Project Work Processes and Organizations: Toward a Micro-Contingency Theory of Organizational Design. *Management Science* 45(11): 1479-1495.

March, J.G. & Simon, H.A. 1958. *Organizations*. New York: Wiley.

McFarlane, D., Sarma, S., Chirn, J.L., Wong, C.Y. & Ashton, K. June 2003. Auto ID systems and intelligent manufacturing control, *Engineering Applications of Artificial Intelligence* 16(4): 365-376.

Muller, J. 9/26/2012. With Driverless Cars, Once Again It Is California Leading The Way. *Forbes*; <http://www.forbes.com/sites/joannmuller/2012/09/26/with-driverless-cars-once-again-it-is-california-leading-the-way/>; retrieved 09/19/2013.

NAVAIR News. November 2011. *Navy completes UAS Common Control System demo*. Available at: <http://www.navair.navy.mil/index.cfm?fuseaction=home.NAVAIRNewsStory&id=4818>

Nissen, M.E. 2013. *Introduction to Command and Control* Naval Postgraduate School CC3000 course, Monterey, CA.

Nissen, M.E. & Buettner, R. June 2004. Computational Experimentation: Bridging the Chasm between Laboratory and Field Research in C2. *Proceedings Command and Control Research and Technology Symposium*, San Diego, CA.

Nissen, M.E. & Place, W.D. June 2013. Archetypical C2 Organization Design for Ever Increasing Technologic Autonomy: An Unmanned Aircraft System Illustration. *Proceedings International Command & Control Research & Technology Symposium*, Alexandria, VA.

Nissen, M.E. & Place, W.D. December 2014, Computational Experimentation to Understand C2 for Teams of Autonomous Systems and People. *Proceedings International Command & Control Research & Technology Symposium*, Alexandria, VA.

Oh, R.P.T, Sanchez, S., Lucas, T., Wan, H. & Nissen, M.E. 2010. Efficient Experimental Design Tools for Exploring Large Simulation Models. *Computational and Mathematical Organization Theory* 15(3): 237-257; DOI 10.1007/s10588-009-9059-1.

Roff, H. August 2015, *The International-Relations Argument Against Killer Robots*. Available at: <http://www.defenseone.com/ideas/2015/08/international-relations-argument-against-killer-robots/119275/>

Selby, G. 2015, Robots Learn to Perform Tasks by “Watchin” Videos. Available at <http://science.dodlive.mil/2015/08/22/robots-learn-to-perform-tasks-by-watching-videos/>.

Simon, H.A. 1991. *Bounded Rationality and Organizational Learning*. Cambridge, MA: MIT Press.

Strauss, A.L. & Corbin, J. 1998. *Basics of Qualitative Research – Techniques and Procedures for Developing Grounded Theory* (2nd Edition) Sage Publications: London.

Thompson, J.D. 1967. *Organizations in Action*. New York, McGraw-Hill.

Thomsen, J. 1998. The Virtual Team Alliance (VTA): Modeling the Effects of Goal Incongruity in Semi-Routine, Fast-Paced Project Organizations. Stanford Univ.

Warne, D. 2015. *Pulse-coupled Neurral Network Performance for Real-time Identification of Vegetation During Forced Landing*. ANZIAM Journal, 5

Work, R. June 2015. Deputy Secretary of Defense Speech at the China Aerospace Studies Institute. As Delivered by Deputy Secretary of Defense Bob Work, RAND Corporation, Arlington, VA. Available at: <http://www.defense.gov/News/Speeches/Speech-View/Article/606683>

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California 93943
3. Research Sponsored Programs Office, Code 41
Naval Postgraduate School
Monterey, CA 93943
4. Dr. Raymond Buettner
CRUSER Director
Naval Postgraduate School
Monterey, CA 93943
5. Dr. Dan Boger
IS Department Chair
Naval Postgraduate School
Monterey, CA 93943