

19th ICCRTS

Toward Computational Modeling of C2 for Teams of Autonomous Systems and People

Autonomy Track – Paper 116

Dr. Mark E. Nissen

US Naval Postgraduate School
1411 Cunningham Road, Monterey CA 93943

W. David Place

US Naval Postgraduate School
Research Associate
&
US Navy - Commander Third Fleet staff
San Diego, CA

Point of Contact: Dr. Mark E. Nissen

US Naval Postgraduate School
MNissen[at]nps.edu

Toward Computational Modeling of C2 for Teams of Autonomous Systems and People

Abstract

The technological capabilities of autonomous systems (AS) continue to accelerate. Although AS are *replacing* people in many skilled mission domains and demanding environmental circumstances, people and machines have complementary capabilities, 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 we identify a plethora of open, command and control (C2) research, policy and decision making questions. Computational modeling and simulation offer unmatched yet largely unexplored potential to address C2 questions along these lines. The central problem is, this kind of C2 organization modeling and simulation capability has yet to be developed and demonstrated in the TASP domain. This is where our ongoing research project begins to make an important contribution. In this article, we motivate and introduce such TASP research, and we provide an overview of the computational environment used to model and simulate TASP C2 organizations and phenomena. We follow in turn with an approach to characterizing a matrix of diverse TASP C2 contexts, as well as a strategy for specifying, tailoring and using this computational environment to conduct experiments to examine such contexts. We conclude then by summarizing our agenda for continued research along these lines.

Introduction

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 technological 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 aerial vehicles 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 technological 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 technological maturity of many AS in operation 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 technological automation of skilled work by machines (e.g., numerical control machining) or employment of computer tools by skilled people (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).

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, for which 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 integration of people and AS into combined units (e.g., integrated squadrons; see C3F, 2013)? 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 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.

Computational modeling and simulation offer an unmatched yet largely unexplored potential to address C2 questions along these lines. If computational models can be 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 could become virtual prototype C2 organizations and 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, technological 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).

The central problem is, this kind of C2 organization modeling and simulation capability has yet to be developed and demonstrated in the TASP domain. Notwithstanding current, lower level work addressing fatigue and like issues affecting individual UxS operators (e.g., see Yang et al., 2012), the higher level C2 modeling and simulation capability envisioned here remains absent.

This is where our ongoing research project begins to make an important contribution. Building upon a half century of research and practice in modeling and simulation in general (e.g., see Forrester, 1961; Law & Kelton, 1991), and two decades of *organization* modeling and simulation work in particular (e.g., see Carley & Prietula, 1994), we have access to computational modeling and simulation technology representing the current state of the art (i.e., VDT [Virtual Design Team]; see Levitt et al., 1999). Such technology leverages well-understood organization micro theories and behaviors that emerge through agent-based interaction (e.g., see Jin & Levitt, 1996). Agent-based organization models developed through this technology have also been validated dozens of times to represent faithfully the structure, behavior and performance of counterpart real-world organizations (e.g., see Levitt, 2004). Plus, we have adapted the same computational modeling and simulation technology over several years to the military domain (i.e., POWER [Project, Organization and Work for edge research]; see Nissen, 2007) to examine joint task forces (e.g., see Looney & Nissen, 2006), distributed operations (e.g., see Oros & Nissen, 2010), computer network operations (e.g., see Koons et al., 2008), and other missions that reflect increasingly common joint and coalition endeavors (e.g., see Gateau et al., 2007).

In the balance of this article, we first provide an overview of the POWER computational environment, which we follow with our approach to characterizing a matrix of diverse TASP C2 contexts, with subsequent discussion of our approach to specifying, tailoring and using POWER to conduct computational experiments to examine such contexts. We conclude by summarizing our agenda for continued research along these lines.

POWER Computational Environment

This section draws heavily from Gateau et al. (2007) to provide an overview of the POWER computational environment. POWER builds upon the planned accumulation of collaborative research over almost 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 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 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 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 Gateau et al. (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 almost two decades, by a team of more than 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.




Figure 2 POWER Model Screenshot

As an example, Figure 2 depicts a screenshot of the POWER computational environment that was used to model a US Military joint task force (JTF) at a relatively high level (e.g., see Gateau et al., 2007). The organization structure is represented by the light (green) person icons at the top of the figure. These correspond to the top three hierarchical levels of the JTF. There are clearly many levels below these that remain unshown in this abstracted model. The task structure is represented by light (yellow) rectangle icons, which are interconnected by dark (black) precedence, medium (red) feedback and other links. The dark (blue) links interconnect organization actors with their tasks (i.e., depicting job assignments), and the medium (purple) trapezoid box at the top represents the set of standing meetings (e.g., Commander’s Brief) that occur routinely. Similarly colored (purple) links indicate which actors are required to participate in such meetings. The interested reader can peruse several articles for details (e.g., see Nissen, 2007; Looney & Nissen, 2006; Oros & Nissen, 2010; Koons et al., 2008; Gateau et al., 2007).

Behind this graphical interface lies the sophisticated modeling and simulation facility of POWER, complete with many dozens of model parameters that can be set to specify a diversity of different C2 organizations and environments. Clearly our TASP C2 model will look different than the JTF representation depicted in the screenshot, and a major aspect of our modeling approach will entail determining how to specify such model in terms of the organization and task structures; their associated links; precedence, feedback, job-assignment and meeting links;

and the many model parameters required to faithfully represent the structure and behavior of TASP C2 organizations and environments in the field.

Modeling Approach

Our focal AS domain in this study centers on the use of multiple unmanned aerial vehicles (UAVs) in an operational military context; that is, a group of UAVs is employed in a potentially hostile environment. More specifically, we focus on UAVs employed onboard one or more ships underway at sea, and we utilize a two dimensional framework to examine a range of increasingly complex employment characteristics in terms of C2. On the one dimension we account for the technological sophistication of the UAVs (Degree 0 – 4); on the other we account for the interdependence between multiple aircraft in concurrent operation (Pooled, Sequential, Reciprocal, Integrated), including both unmanned-only and integrated manned-unmanned missions. The five sophistication degrees derive from the domain of autonomous automobiles and are set by the National Highway Traffic Safety Administration (Fisher, 2013). Here we outline first the five degrees for autonomous automobiles, then we map such degrees to the UAV domain.

Briefly, in the autonomous automobile domain, Degree 0 corresponds to no autonomy; the car must be controlled continuously by a person in the driver's seat. Degree 1 corresponds to incorporation of standard safety features (e.g., antilock brake system [ABS], electronic stability control [ESS], adaptive cruise control [ACC]) that assist the driver with one specific aspect of controlling a vehicle. Degree 2 corresponds to two or more Degree 1 capabilities (e.g., automatic lane centering and adaptive cruise control) that integrate to enable a car to drive itself to a limited extent (e.g., within one particular lane of a specific road; person in driver's seat ready to take control at any time). Degree 3 corresponds to incorporation of an autopilot, which enables the car to change lanes and roads to reach a predetermined destination, but the driver must stay engaged and ready to resume control if the car gets confused or into a situation beyond its capability. Degree 4 corresponds to a car that can start and complete an entire trip without human engagement (e.g., no driver or passengers; no one in driver's seat).

Mapping this loosely to the UAV domain, Degree 0 could describe a (manned) aircraft that must be controlled continuously by a person in the cockpit; this represents a relatively direct mapping from the automobile domain to its UAV counterpart. Degree 1 could describe an aircraft (esp. UAV) that can be controlled continuously by a person via remote control (no one in the cockpit); this cross-domain mapping is less direct than for Degree 0, but it incorporates a key difference between cars and UAVs: remotely controlled cars are hobbyists' toys, whereas remotely piloted UAVs fly serious missions. Degree 2 could describe a manned aircraft on autopilot, but for our present purpose we characterize it instead in terms of a UAV that can fly between two, pre-established waypoints but with a human controller ready to take

control when the waypoint is reached; this returns us to a relatively direct mapping from one domain to the other. Degree 3 could describe a UAV that can navigate to multiple waypoints and destinations with no remote control while in flight, but such UAV would require human control to take off, land, recover from unexpected events, and complete certain missions (e.g., changing surveillance modes, releasing ordnance, recognizing and pursuing targets of opportunity). Degree 4 could describe a UAV capable of taking off, landing, handling the unexpected, and completing all of its assigned missions with no human control or intervention. Table 2 summarizes the cross-domain autonomy degree mapping between the automobile and UAV.

Table 2 Cross-Domain Autonomy Degree Mapping

Degree	Automobile	UAV
0	No autonomy; continuous human control	Manned aircraft; continuous human control
1	Safety features (ABS, ESS, ACC)	Remotely controlled aircraft (UAV)
2	Limited autonomous driving (lane control)	Autonomous flight between waypoints
3	Autopilot (lane & road changes)	Autonomous navigation (multiple waypoints)
4	Full autonomy; human driver not required	Full autonomy; human controller not required

The interdependence dimension derives from Organization Theory (Thompson, 1967). It characterizes the intensity of interactions and behaviors within an organization. At its most basic, Pooled interdependence describes how different units of an organization (e.g., different departments, groups, functions) can each contribute to the overall operation and success of the organization but without direct interaction with one another. An organization’s legal department and its building maintenance unit reflect pooled interdependence as such; they both contribute to the same organization’s overall operation and success, but the legal and maintenance units do not interact with one another commonly. Coordination between units characterized by pooled interdependence is minimal and accomplished through rules and standards generally, for each unit operates independently.

Sequential interdependence subsumes its pooled counterpart but incorporates the additional interactions associated with one unit in the organization producing outputs necessary for subsequent performance by another unit. An organization’s engineering and manufacturing units reflect sequential interdependence as such; the designs developed within the engineering unit are used as inputs to the products built within the manufacturing unit. Coordination between units characterized by sequential interdependence is more intensive and accomplished via plans and schedules generally.

Reciprocal interdependence subsumes its pooled and sequential counterparts but incorporates the additional interactions associated with two units working simultaneously on a

common task. A surgeon and nurse operating on a patient reflect reciprocal interdependence as such; the surgeon cannot perform certain tasks until the nurse has performed his or hers, and vice versa, over time, nor can either surgeon or nurse anticipate all possible outcomes or issues that might emerge through surgery (e.g., they must observe and communicate together, and they must react and adjust jointly as the surgery progresses). Coordination between units characterized by reciprocal interdependence is highly intensive and accomplished via recurring feedback and mutual adjustment.

In the UAV domain, pooled interdependence could refer to two or more, different aircraft—manned or unmanned—that contribute to the overall operation and success of the organization but without direct interaction with one another. Say that two different aircraft perform surveillance missions in different geographical areas. The surveillance from both aircraft is useful to the organization, but neither aircraft interacts with the other. Coordination can be via specific deconfliction rules, for instance, that prohibit two aircraft from flying in the same airspace at the same time.

Sequential interdependence could refer to two or more, different aircraft that share pooled interdependence but also depend upon one another over time. Say that one aircraft performs a surveillance mission and provides targeting information for a different type of aircraft to attack. Coordination can be via air plans, for instance, that schedule the second aircraft to fly after receiving useful targeting information from the first one.

Reciprocal interdependence could refer to two or more, different aircraft that share pooled and sequential interdependence but must also work simultaneously on a common task. Say that two aircraft are required to defend one another if either is attacked. Coordination requires frequent communication between the aircraft, for instance, and both must adjust their actions depending upon circumstances.

Table 3 summarizes this interdependence scheme for the organization and UAV domains. Notice that we include the Integrated interdependence type also. This is not part of the organization scheme noted above, but we incorporate it to characterize two different *organizations* that work together in manners reflecting reciprocal interdependence. Hence, beyond having two different units *within* the same organization performing reciprocally, such units must do so *across* different organizations, for example in a joint project where neither organization is solely “in charge” of the whole effort; many strategic partnerships, joint spinoffs and complex endeavors reflect this property (e.g., see Alberts & Hayes, 2006). In the UAV domain, we use this interdependence type to characterize reciprocally interdependent missions with both manned and unmanned aircraft “organizations” flying and working together toward a common objective.

Table 3 Interdependence Scheme

Interdependence Type	Organization	UAV
Pooled	Minimal interaction; coordination via rules & standards	Aircraft performing surveillance missions in different geographic areas
Sequential	Outputs from one organization unit are inputs to another; coordination via plans & schedules	Surveillance from one aircraft provides targeting information for another
Reciprocal	Two or more units perform a common task; coordination via feedback & mutual adjustment	Two aircraft defend one another if either is attacked
Integrated	Two or more different organizations perform a common task reflecting reciprocal interdependence.	Manned and unmanned aircraft fly together and defend one another.

With these two dimensions, we can consider—in a systematic and orderly manner—a 5x4 matrix of increasing complex TASP C2 contexts. We summarize this context matrix in Table 4. 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 4 UAVs reflecting both reciprocal interdependence among themselves and integrated interdependence with their manned aircraft counterparts (i.e., labeled “D4I” in the table). Each of the key intermediate conditions (i.e., Degree 1 to Degree 4 sophistication, across all four interdependence conditions) is examined systematically also for completeness. This matrix specifies the set of computational experiments to be conducted.

Table 4 Matrix of TASP C2 Contexts to Examine

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

POWER Specification, Tailoring and Use

As noted above, POWER is designed and validated to represent and simulate the structures and behaviors of organizations in a manner that supports computational experiments. Such design

and validation focus on *people* in the organization that use many different kinds of tools, machines and other technologies to perform work. To the extent that our TASP C2 context centers on people using aircraft, communication and other technologies to accomplish work, the POWER computational environment should serve us very well, for as noted in the introduction, we have adapted it for and validated it in the military domain previously (e.g., see Nissen, 2007; Looney & Nissen, 2006; Oros & Nissen, 2010; Koons et al., 2008; Gateau et al., 2007).

However, a central aspect of our TASP C2 context involves robots, unmanned vehicles and other autonomous systems (esp. UAVs), which by definition operate *without people* controlling them (esp. Degree 4 AS). The POWER model structures and behaviors may not be specified to represent and simulate such autonomous systems well. We will need to understand both the common and unique aspects of AS with respect to their human counterparts, and we will need to assess POWER's extant capability to represent and simulate them faithfully. This represents a key aspect of our ongoing research.

Alternatively, many of the conditions summarized in our TASP C2 context matrix above conform well to people using technological tools, hence POWER should work well with them in its current state. For instance, all of the conditions reflecting Degree 0 and 1 sophistication (i.e., across all interdependence cases) appear to be well within extant POWER capability, and one can argue that those reflecting Degree 2 and 3 sophistication are within such capability too, for humans remain in charge of machines and are ready to retake control at any time. This is not much different than a human operating a machine that is capable of performing a limited set of actions on its own but that requires human input and attention to perform the complete set.

An air traffic controller (ATC), as one example, uses sophisticated radar, computer and communication technologies to keep track of and manage myriad aircraft flying through their assigned regions of airspace. Although many such technologies can operate independently (e.g., automatic radar position tracking)—and the aircraft themselves are capable of flying without ATC or pilot input—the ATC remains in charge of the airspace and is ready to control the aircraft's position at any time (esp. in case of potential collision or emergency). Another example, albeit somewhat trivial, pertains to the exceedingly common case of a person using a washing machine to clean laundry. Once loaded with laundry and detergent, and set for the desired water level and temperature, the washing machine can complete the cleaning cycle without human intervention. Nonetheless, few washing machines can load or unload themselves, and the human must at least monitor the machine in case it gets off balance or manifests some other issue. POWER can model these and like cases well in its present condition.

As a purposefully naïve, initial position—akin to a null hypothesis if you will—we can *make a big assumption* and consider Degree 4 UAVs and other AS to behave in manners that

are consistent with the behaviors of their human counterparts in our C2 organization context. On the one hand, this is not a great leap of faith, for most AS are designed to emulate human behaviors, and the more sophisticated the AS, the more closely its behavior mirrors that of human counterparts. On the other hand, however, people and machines possess different characteristics and capabilities (e.g., machines excel at consistency, memory, processing power, endurance; people excel at judgment, innovation, adaptation and working with uncertainty), and understanding their relative behaviors in the C2 organization context demonstrates both the novelty and potential of our present line of research.

Conclusion

The technological capabilities of autonomous systems (AS) continue to accelerate. Although AS are *replacing* people in many skilled mission domains and demanding environmental circumstances, people and machines have complementary capabilities, and *integrated* performance by AS and people *working together* can be superior. We refer to this increasingly important phenomenon as Teams of Autonomous Systems and People (TASP), and we identify a plethora of open, C2 research, policy and decision making questions.

Computational modeling and simulation offer an unmatched yet largely unexplored potential to address C2 questions along these lines. The central problem is, this kind of C2 organization modeling and simulation capability has yet to be developed and demonstrated in the TASP domain. This is where our ongoing research project begins to make an important contribution.

In this article, we motivate and introduce such TASP research, and we provide an overview of the POWER computational environment used to model and simulate TASP C2 organizations and phenomena. POWER offers a powerful and unique capability to represent and simulate organizations and associated phenomena, particularly as it has been validated many times across a variety of domains, including military operations.

We follow in turn with our approach to characterizing a matrix of diverse TASP C2 contexts. Drawing from the domain of autonomous automobiles, we identify and outline a scheme for characterizing and operationalizing five increasing degrees of autonomy, ranging from *none* to *full*. We then map such scheme to the UAV domain and establish a useful dimension for characterization and analysis in terms of AS sophistication.

Similarly, drawing from the domain of Organization Theory, we identify and outline a scheme for characterizing and operationalizing four increasing types of interdependence between organization units, ranging from *pooled* to *integrated*. We then map such scheme to the UAV domain and establish another useful dimension for characterization and analysis. The resulting 5x4 matrix characterizes a broad set of diverse TASP C2 contexts, which specifies the set of computational experiments to be conducted.

This matrix of TASP C2 contexts represents a contribution of new knowledge. As noted above, it serves to specify and can drive our subsequent computational experiment. However, it may also prove to be useful to other researchers, leaders, managers, practitioners and policy makers. Each cell in this context matrix represents a different set of conditions associated with UAV planning and operation, and one can envision a different set of techniques, tactics and

procedures corresponding to each, for instance. This matrix also outlines a wide range of technical and operational conditions to be considered, which researchers can leverage to identify high value targets of future research, for instance.

Additionally, we discuss issues with and our approach to specifying, tailoring and using this POWER computational environment to conduct experiments to examine the matrix of TASP C2 contexts. Although POWER appears to be quite capable in its current state to represent and simulate many of the conditions within the context matrix, some of the more extreme conditions (esp. Degree 4 AS, integrated interdependence) may prove to be a challenge. A major thrust of ongoing research following the present study is examining exactly such issues in terms of representation and simulation. Our purposefully naïve, initial position—a null hypothesis if you will—is that fully autonomous systems behave like people do, but we realize that considerable research is required to validate, refute or refine said position.

Indeed, working with experts in the UAV domain, our ongoing research is examining systematically how various AS do and do not behave like their human counterparts do. Because so many diverse AS—like their human counterparts—can potentially behave differently, uniquely and idiosyncratically across various domains, we focus our investigation initially on UAVs employed by military organizations from ships underway at sea, and we look first at surveillance missions.

Nonetheless, understanding this domain well—and establishing the capability for POWER to represent and simulate the corresponding structures and behaviors faithfully—will produce new knowledge and guide C2 organization leaders, managers, practitioners and policy makers in their planning, operation, use, specification and acquisition of UAVs and like AS. This is in addition of course to guiding other C2 and organization researchers as well. It is exciting to be conducting cutting-edge research at this unique point in the advance of technology and organization pertaining to AS. We remain enthused by our own research, and we encourage others to contribute in complementary ways.

References

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.

C3F. 30 April 2013. Commander Third Fleet TACMEMO 2-01.1-13, Ship-Based Unmanned Aircraft Systems (UAS) Employment Tactics.

Carley, K.M. & Prietula, M.J. 1994. *Computational Organization Theory*. Hillsdale, NJ, Lawrence Erlbaum Associates.

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

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.

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

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 code19991004 162; retrieved 11/05/2012.

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

Fisher, A. 2013. Robots Can Already Outdrive Humans: Now Everyone Needs to Get Out of Their Way. *Popular Science* 283(4): 54-62.

Forrester, J.W. 1961. *Industrial Dynamics*. Pegasus Communications.

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.

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 Organization Theory 2(3): 171-195.

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.

Koons, J.L., Bekatoros, N. & Nissen, M. June 2008. C2 for Computer Networked Operations: Using Computational Experimentation to Identify Effects on Performance in Organizational Configurations within the Larger Network-Centric Environment. *Proceedings International Command & Control Research & Technology Symposium*, Bellevue, WA.

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.

Looney, J.P. & Nissen, M.E. June 2006. Computational Modeling and Analysis of Networked Organizational Planning in a Coalition Maritime Strike Environment. *Proceedings Command & Control Research & Technology Symposium*, San Diego, CA; Winner – Best Student Paper Award.

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.

Nissen, M.E. 2007. Computational Experimentation on New Organizational Forms: Exploring Behavior and Performance of Edge Organizations. *Computational and Mathematical Organization Theory* 13(3): 203-240.

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 Technological Autonomy: An Unmanned Aircraft System Illustration. *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.

Oros, C. & Nissen, M.E. 2010. Designing Complex Organizations Computationally. in: M. Wang and Z. Sun (Eds.), *Handbook of Research on Complex Dynamic Process Management: Techniques for Adaptability in Turbulent Environments*. Hershey, PA: IGI Global: 573-598.

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.

Yang, J., Kapolka, M. & Chung, T. 2012. Autonomy balancing in a manned-unmanned teaming (MUT) swarm attack. The 1st International Conference on Robot Intelligence Technology and Application (RiTA), Gwangju, Korea.