

Computational Experimentation with the Virtual Design Team: Bridging the Chasm between Laboratory and Field Research in C2¹

Mark E. Nissen²
Raymond R. Buettner

Naval Postgraduate School
555 Dyer Road, Code GB/Ni, Monterey, CA 93943-5000
Tel: 831-656-3570; fax: 831-656-3407; e-mail: MNissen@nps.navy.mil

Stanford University
394 Terman, Stanford, CA 93943-5000
Tel: 831-656-3387; fax: 831-656-3407; e-mail: Buettner@nps.navy.mil

Student Paper

Abstract

A chasm exists between laboratory and field methods in C2 research. These methods are complementary but used rarely in combination. This expository article describes a research approach that bridges such chasm: computational experimentation. Computational experimentation mitigates the weakness of both laboratory and field research, yet it has its own limitations and appears suited best as a complement and not a replacement. To illustrate the power and potential of computational experimentation, we describe an implemented agent-based modeling environment called VDT. VDT benefits from accumulated research over two decades and extensive external validation. We employ this modeling environment to represent and emulate the behavior of a high-level C2 organization. Using a full-factorial experimental design, we illustrate computational experimentation through controlled manipulation of key factors associated with organizational and technological design (i.e., bureaucracy level, coordination load, knowledge inventory). This illustration includes discussion of rich operationalized constructs used to characterize a diversity of C2 organizations, task environments and performance measures. The experimental results highlight complex interactions between design factors, and they suggest fundamental tension and decision tradeoffs between important performance measures such as mission duration and risk. The article closes with key conclusions, implications for C2 in practice today, and suggestions for future research.

Introduction

Command and control (C2) represents a complex system (Jackson and Keyes 1984) that often involves large and distributed organizations, diverse and sophisticated technologies, and a mixture of novice and highly trained personnel, operating in hazardous and equivocal

¹ Forthcoming: *Proceedings* Command & Control Research & Technology Symposium, San Diego, CA (June 2004).

² This research is sponsored in part by the Office of Naval Research, Young Investigator Program: N0001401WR20304.

environments usually under the pressure of time constraints. As such, a huge number of diverse factors combine in various ways and at various times to influence the performance of C2 processes. Such number and diversity make it very difficult to attribute causality and identify which specific combinations of organizational structures, work processes, technologies and personnel contribute to relatively better and worse performance. Hence trial and error abounds in C2 practice, as does perpetuation of what has appeared to work well—or at least not failed—in the past.

Through research in the C2 domain, new knowledge continues to develop through a variety of methods. For instance, analytical methods such as optimization can be employed quickly and at low cost to combine myriad variables and identify the best mix to pursue some particular objective and satisfy specific constraints. As another instance, laboratory methods such as experimentation can also be employed relatively quickly and at low cost to control factors and manipulate variables to trace causality in C2 processes. This class of analytical and laboratory methods suffers from well-known weaknesses, however (e.g., poor external validity & generalizability). Alternatively, field methods such as case study and action research can be employed to study professional people in operational organizations performing complex work processes at sea and in other hazardous environments. But this class of field methods also suffers from well-known weaknesses (e.g., high cost & time consuming, poor experimental control & internal validity). Even field experiments exhibit great difficulty with control, validity and confounding at present. The relative weaknesses associated with these two classes points to a chasm that exists between laboratory and field methods in C2 research.

This expository article describes a research approach that bridges the chasm between laboratory and field methods in C2: computational experimentation. As a bridge method, it mitigates the weaknesses of research methods in both classes (i.e., laboratory and field) and hence offers relative advantages over either method individually. However, computational experimentation has its own limitations and hence appears suited best to be used to complement laboratory and/or field methods, not replace them. To illustrate the power and potential of computational experimentation, we describe an implemented agent-based modeling environment called the VDT. VDT benefits from accumulated research over two decades and extensive external validation. We employ this modeling environment to represent and emulate the behavior of a high-level C2 organization. Using a full-factorial experimental design, we illustrate computational experimentation through controlled manipulation of key factors associated with organizational and technological design (i.e., bureaucracy level, coordination load, knowledge inventory). The article closes with key conclusions, implications for C2 in practice today, and suggestions for future research.

Computational Experimentation

Throughout the era of modern science a chasm has persisted between laboratory and field research. On one side the laboratory provides unparalleled opportunity for controlled experimentation. Through experimentation the researcher can manipulate only a few variables of interest at a time and can minimize the confounding associated with the myriad factors affecting complex systems and processes in the field (Box et al. 1978, Johnson and Wichern 1992). However, limitations of laboratory experimentation are known well (Campbell and Stanley 1973) and particularly severe in the domain of command and control (C2). In C2 experimentation such limitations center on problems with external validity. Laboratory conditions can seldom replicate the complexity, scope and scale of the physical organizations

and systems of interest for research. Experiments also include problems with generalizability. Many experiments utilize samples of convenience (esp. university students) instead of working professionals. This practice calls into question how closely the associated experimental results are representative of C2 behavior in operational organizations. These same concerns also pertain to analytical methods (e.g., mathematical analysis, optimization; see Chiang 1984, Lapin 1985). Most such methods use theoretical concepts as variables, not operationalized constructs, and of course analytical models do not involve real people, systems and organizations.

On the other side field research provides unparalleled opportunity for realism (Denzin and Lincoln 1994). The researcher in the field can study full-scale artifacts in operational environments (Yin 1994) and can minimize the abstraction away from working people, systems and organizations (Glaser and Strauss 1967). However, limitations of field research are also known well (Campbell and Stanley 1973) and particularly severe in the C2 domain. In C2 field research such limitations center on problems with internal validity. Field research affords little opportunity for controlled experimentation (cf. Cook and Campbell 1979). Also, confounding results often from the myriad influences on complex systems and organizations that cannot be isolated in the field. This practice makes it difficult to identify and trace the causes of differential behaviors—better as well as worse—in C2.

As implied by the name, computational experiments are conducted via computer simulation. As such they offer all of the cost and time advantages of computational analysis. But computational experiments go beyond most simulations. Rigorous experimental designs are employed to capture the benefits of laboratory experimentation. The variables affecting physical systems and organizations in the field can be isolated and examined under controlled conditions. This also addresses the internal validity and confounding limitations of field research. Yet computational experiments can be conducted at a fraction of the cost and time required to set up and run experiments with human subjects in the laboratory. Further, through external validation, computational models can demonstrate fidelity emulation of the key qualitative and quantitative behaviors of the physical systems and organizations they represent. This addresses the problems with external validity and generalizability noted above.

It is important to note, computational modeling and simulation are not new techniques for the study of C2. For instance, the Adaptive Architectures for Command and Control (A2C2) team (see Diedrich et al. 2003, Kleinman et al. 2003) employs a “Model Driven Experimentation Paradigm” (Handley 1999). But this method sits squarely within the class of analytical and laboratory methods noted above, using analytical models to guide laboratory experimentation. Carley (1999) uses computational methods for hypothesis generation, as another instance.

Figure 1 illustrates the essential elements of computational experimentation as a research method. The top of the figure includes a shape to depict the bridge metaphor associated with this method, as it spans a wide gap between laboratory and field methods. From the left side of this “bridge,” two arrows represent inputs to describe the behaviors of computational models. Organization theory, which is predicated upon many thousands of studies over the last half century, provides the basis for most such behaviors. Behaviors pertaining to organizational factors such as centralization, division of labor, task interdependence, function, coordination, formalization, technology and information processing are captured from organization theory. Where extant theory does not address well a behavior of interest (e.g., knowledge flows), ethnographic and like immersive field studies (Bernard 1998) are conducted to understand the associated organizational behaviors. Because organization theory is general, and not based on any single organization, the associated behaviors have broad applicability across organizations in

practice. This provides for the generalizability attainable through the method of computational experimentation.

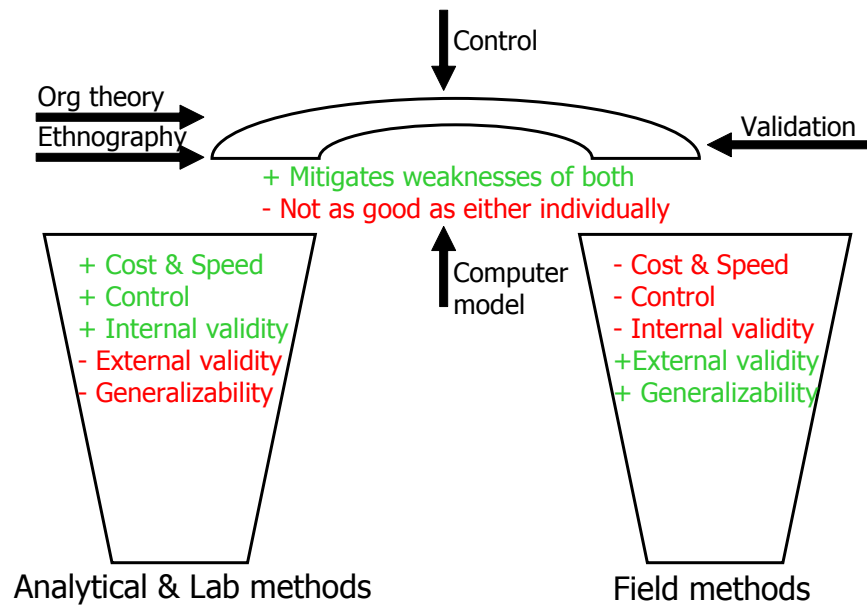


Figure 1 Bridge Method

From the bottom of the “bridge,” an arrow represents the use of computer models to represent organizations and emulate their key behaviors. Some variety exists in terms of specific implementations, but most computer models adhere to standards, norms and conventions associated with the field of Computational Organization Theory (COT; see Carley and Prietula 1994). The central goal is to develop computer models that emulate the key behaviors of organizations and to use such models to examine alternate methods of organization and coordination. As such COT shares a focus on many factors of importance in command and control.

From the right side of the “bridge” in the figure, one arrow represents a requirement in our approach for model validation. Through validation, the organizational behaviors emulated by computer models are examined and compared with those of operational organizations in the field. We view this as an essential step, for it provides confidence that the behaviors emulated by computer model have sufficient fidelity to mirror faithfully the behaviors of the operational organizations they represent. This provides for the external validity attainable through the method of computational experimentation. It is important to note, not all COT models are subjected to such validation. Many researchers use computational models to conduct theorem proving studies, which are valuable in their own right to demonstrate various aspects of organization theory. But without validation, such researchers have difficulty making claims that such theory mirrors the behavior of organizations in the field. Hence validation represents an important characteristic to distinguish computational experimentation as the research method described specifically in this article from COT in general.

Finally, from the top of the “bridge,” an arrow represents the use of experimental controls in research. Following the same rich set of experimental designs available to laboratory researchers (e.g., full-factorial, Latin Squares, blocking with replication), computational experimentation as a research method can be used to control for myriad factors and manipulate just one or a few variables at a time to examine causality. Further, the same experimental design and setup can be replicated any number of times, for instance using Monte Carlo techniques or other computational approaches to introduce variation. This provides for the internal validity attainable through the method of computational experimentation. Combining these “bridge” inputs together—organization theory and ethnography, computer models, validation and control—the method of computational experimentation can be understood in terms of, and indeed inherits, the various properties of its constituent elements.

Figure 1 also illustrates the bridging nature of computational experimentation as a research method. On the left side we depict analytical and laboratory methods and summarize their key advantages (e.g., low-cost & fast studies, good experimental control & internal validity) and disadvantages (e.g., poor external validity & generalizability). On the right side we depict field methods in similar fashion to summarize their key advantages (e.g., good external validity and generalizability) and disadvantages (e.g., high cost & time consuming, poor experimental control & internal validity). Notice from their relative advantages and disadvantages how the two classes of research methods complement one another. Field methods are strong in the areas where analytical and laboratory methods are weak, and vice versa. As an alternate research method, computational experimentation mitigates weakness of both classes. For instance, it enables good experimental control and internal validity as in laboratory methods, yet also promotes good generalizability and external validity as in field methods. Nonetheless, every research method is flawed in some respects. In our present case, when used in isolation, computational experimentation is not as good as either method at its best. For instance, because computational experimentation uses computer models of people in organizations instead of real people, it is weaker in this respect than laboratory experimentation is. This same use of computer models instead of real people also makes computational experimentation weaker than field methods are. This is why we describe computational experimentation as a *bridge method*: it bridges the chasm between experimental and field research methods, but it serves best to complement, not replace, such methods.

Agent-Based Modeling Environment

In this section, we build upon current advances in VDT research—which represents a branch of computational organization theory—to describe the agent-based modeling environment used here for computational experimentation. Drawing heavily from Nissen and Levitt (2004), we first summarize the stream of research associated with VDT and then describe its modeling environment.

Virtual Design Team Research

The Virtual Design Team (VDT) Research Program (VDT 2004) reflects the planned accumulation of collaborative research over two decades to develop rich theory-based models of organizational processes. Using an agent-based representation (Cohen 1992, Kunz et al. 1998), micro-level organizational 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 representational fidelity and have shown

how the emulated behaviors of VDT computational models correspond closely with a diversity of enterprise processes in practice.

The VDT research program continues with the goal of developing new micro-organization theory and of 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 alternate prototype systems. Clearly this represents a significant challenge in the domain of organizations. Micro-theory and analysis tools for designing bridges and airplanes rest on well-understood principles of physics (e.g., involving continuous numerical variables, describing materials whose properties are relatively easy to measure and calibrate), and analysis of such physical systems yields easily to differential equations and precise numerical computing.

In contrast, theories describing the behavior of organizations are characterized by nominal and ordinal variables, with poor measurement reproducibility, and verbal descriptions reflecting significant ambiguity. Unlike the mathematically representable and analyzable micro-behaviors of physical systems, the dynamics of organizations are influenced by a variety of social, technical and cultural factors, are difficult to verify experimentally, and are not as amenable to numerical representation, mathematical analysis or precise measurement. Moreover, quite distinct from physical systems, people and social interactions—not molecules and physical forces—drive the behavior of organizations. Hence such behaviors are fundamentally non-deterministic and difficult to predict at the individual level. Thus, people, organizations and business processes are qualitatively different than bridges, semiconductors and airplanes, and it is irrational to expect the former to ever be as understandable, analyzable or predictable as the latter. This represents a fundamental limitation of the approach.

Within the constraints of this limitation, however, we can still take great strides beyond relying upon informal and ambiguous, natural language textual description of organizational behavior (e.g., the bulk of extant theory). For instance, the domain of organization theory is imbued with a rich, time-tested collection of micro-theories that lend themselves to qualitative 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 from this theory base, 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 emulate the dynamics of organizational behaviors.

Even though the representation is qualitative (e.g., lacking the precision offered by numerical models), through commitment to computational modeling, it becomes semi-formal (e.g., different people viewing the model can agree on what it describes), reliable (e.g., the same sets of organizational conditions and environmental factors generate the same sets of behaviors), and explicit (e.g., much ambiguity inherent in natural language is obviated). Particularly when used *in conjunction with* the descriptive natural language theory of our extant literature, this represents a substantial advance. Further, once a model has been validated to emulate accurately the qualitative behaviors of the field organization it represents, it can be used to examine a multitude of cases (e.g., many more and diverse than observable in practice) under controlled conditions (e.g., repeating the same events multiple times, manipulating only one or a few variables at a time through repeated trials, stopping the action for interpretation). This alone offers great promise in terms of theory development and testing.

Additionally, although organizations are inherently less understandable, analyzable and predictable than physical systems are, and the behavior of people is non-deterministic and difficult to model at the individual level, it is known well that individual differences tend to average out when aggregated cross-sectionally and/or longitudinally. Thus, when modeling aggregations of people in the organizational context (e.g., work groups, departments, 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 unpredictable attention of a worker to any particular activity or event (e.g., new work task, communication, request for assistance) can be modeled—stochastically—to approximate collective behavior. As another instance, specific organizational behaviors can be simulated hundreds of times—such as through Monte Carlo techniques—to gain insight into which results are common and expected versus those that are rare and exceptional.

Of course, applying numerical simulation techniques to organizations is nothing new (e.g., see Law and Kelton 1991). But this approach enables us to *integrate* the kinds of dynamic, qualitative behaviors emulated by symbolic models with quantitative aggregate dynamics generated through discrete-event simulation. It is through such integration of qualitative and quantitative models—bolstered by strong reliance upon well-established theory and commitment to empirical validation—that our approach diverges most from extant research methods and offers new insight into the dynamics of organizational behavior.

VDT Modeling Environment

Here we provide a brief overview of the VDT modeling environment. The development and evolution of VDT has been described in considerable detail elsewhere (e.g., Cohen 1992, Christiansen 1993, Jin and Levitt 1996, Thomsen 1998, Kunz et al. 1998, Levitt et al. 1999, Nogueira 2000, VDT 2004), so we do not repeat such discussion here. The VDT modeling environment has been developed directly from Galbraith's information processing view of organizations. This information processing view has two key implications (Jin and Levitt 1996). 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. In essence this amounts to overlaying the task structure on the organization structure and to developing computational agents with various capabilities to emulate the behaviors of organizational actors performing work.

Figure 2 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 organizational 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 organizational attributes such as skills possessed, level 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.

As also suggested by the figure, 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 also 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. Collective task performance is constrained further by the number of individual actors assigned to each task, the magnitude of the task, and both scheduled (e.g., work breaks, ends of shifts, weekends and holidays) and unscheduled (e.g., awaiting managerial decisions, awaiting work or information inputs from others, performing rework) downtime.

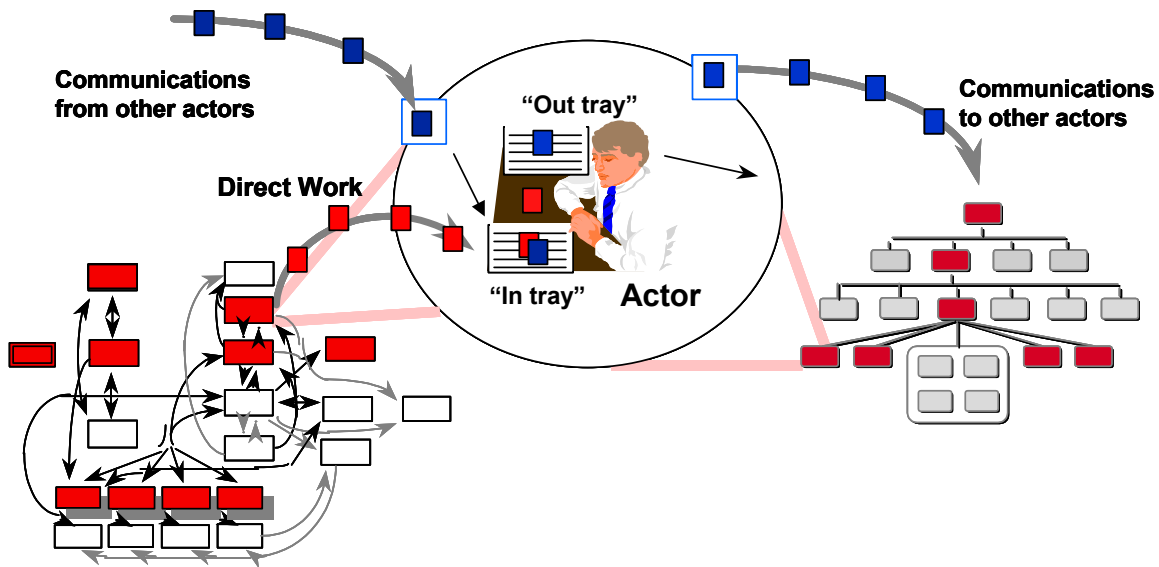


Figure 2 VDT Information Processing View of Knowledge Work

The second implication is computational: both primary work (e.g., planning, design, management) and coordination work (e.g., group tasks, meetings, joint problem solving) are modeled in terms of *work volume*. This construct is used to represent a unit of work (e.g., associated with a task, a meeting, a communication) within the task structure. In addition to symbolic execution of VDT models (e.g., qualitatively assessing skill mismatches, task-concurrency difficulties, decentralization effects) through micro-behaviors derived from organization theory, the discrete-event simulation engine enables (virtual) process performance to be assessed (e.g., quantitatively projecting task duration, cost, rework, process quality).

Clearly quantitative simulation places additional burden on the modeler in terms of validating the representation of a knowledge-work process, which generally requires fieldwork to study an organization in action. The VDT modeling environment benefits from extensive fieldwork in many diverse enterprise domains (e.g., power plant construction and offshore drilling, see Christiansen 1993; aerospace, see Thomsen 1998; software development, see Nogueira 2000; healthcare, see Cheng and Levitt 2001; others). Through the process of “backcasting”—predicting known organizational outcomes using only information that was available at the beginning of a project—VDT models of operational enterprises in practice have

demonstrated dozens of times that emulated organizational behaviors and results correspond qualitatively and quantitatively to their actual counterparts in the field (Kunz et al. 1998).

Viewing VDT as a validated model of project-oriented knowledge work, researchers have begun to use this dynamic modeling environment as a “virtual organizational testbench” to explore a variety of organizational questions, such as effects of distance on performance (Wong and Burton 2000), or to replicate classic empirical findings (Carroll and Burton 2000). Thus the VDT modeling environment has been validated repeatedly and longitudinally as representative of both organization theory and enterprises in practice. This gives us considerable confidence in its results. Moreover, VDT is designed specifically to model the kinds of knowledge work and information processing tasks that comprise the bulk of C2 processes.

VDT Command and Control Model

Here we employ the VDT modeling environment to represent work processes associated with a high-level command and control organization. The organization modeled here is fictitious but representative at a general level of those involved with large-scale C2. VDT is capable of modeling large, complex, operational organizations in great detail, and it has been demonstrated repeatedly to emulate well the associated behaviors of organizations in the field. But using a high-level model as such helps us maintain the focus of this expository article on techniques of VDT modeling and computational experimentation, which represents our primary contribution, and not get lost in the details of the organization itself. We first describe the VDT representation and then illustrate how a full-factorial computational experiment can be performed on it.

VDT C2 Model

Figure 3 presents a screenshot of this high-level VDT C2 model. The model is comprised of four mission tasks (i.e., denoted as lightly colored boxes) divided in to two phases. In the first phase, air and surface missions (i.e., labeled “Air Missions 1” and “Surface Missions,” respectively) are planned for coordinated execution. Upon successful completion of both missions, a milestone marking the beginning of Phase 2 is noted. In the second phase, a different set of air missions is planned for coordinated execution with a ground assault (i.e., labeled “Air Missions 2” and “Ground Missions,” respectively). Mission tasks require resources to perform, demand particular capabilities, and vary in terms of magnitude, complexity and timing.

The coordination links (i.e., denoted by light dashed lines) connecting the coordinated missions denote reciprocal task interdependencies (Thompson 1967), which suggest they must be coordinated closely in both planning and execution. For instance, an air mission such as removing anti-ship assets must be coordinated closely with anti-mining and shipboard artillery surface missions. Likewise, close-air support must be coordinated closely with amphibious ground assault. VDT emulates the added coordination effort associated with such reciprocal task interdependencies. The rework links (i.e., denoted by dark dashed lines) connecting tasks from different mission phases denote sequential task dependencies, which suggest the predecessor missions must be accomplished effectively in order for the successors to perform well. In the case of amphibious assault, for instance, this depends heavily upon success of the anti-mining and artillery operations. To the extent that such predecessor work is not completed or not accomplished effectively, certain aspects may have to be redone to correct any major deficiencies.

The people icons depict organizations and are arranged in terms of the command hierarchy. People icons represent resources, which have particular capabilities, skill levels and

roles. Where a skilled actor’s capability matches that required for a mission task, the resource is likely to perform it competently and within the time required. If the actor has greater or lesser skill, the time required to perform the mission task can be appreciably shorter or longer, and the competency of performance can be notably better or worse, respectively. Where an actor does not possess the required capability at all, the mission task will be in jeopardy. Such relationships are appealing intuitively and reflect well many organizational behaviors.

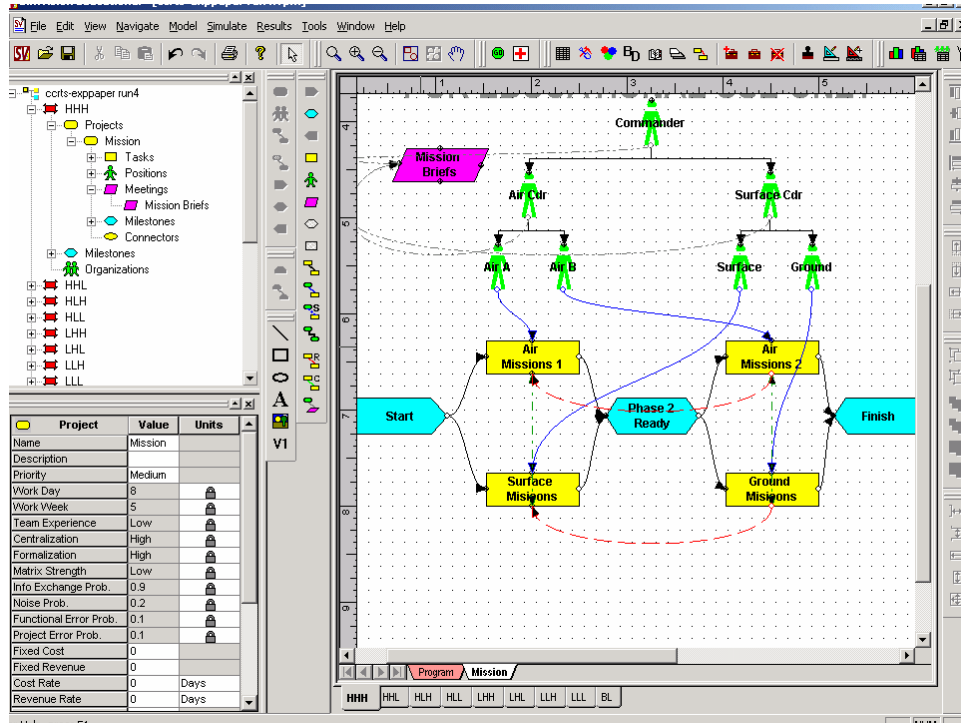


Figure 3 VDT C2 Model Screenshot

A Commander actor sits atop the organization and has two major command organizations reporting to it: 1) an Air Command is responsible for all air missions, and 2) a Surface Command is responsible for both surface and amphibious missions. Reporting to each of these commands is a set of actors with different capabilities. For instance, the icon labeled “Air A” could represent an aviation organization specializing in reconnaissance and strike warfare, and its counterpart labeled “Air B” could specialize instead on aerial attack and support. Likewise, the “Surface” unit could involve ships at sea, and the “Ground” unit could be comprised of expeditionary forces. Notice the VDT representation includes a mission task structure and an organization structure. The assignment links (i.e., delineated by solid lines) denote which organizational actors are responsible for the various mission tasks. Finally, a dark trapezoid box is used to depict recurring meetings that must be attended by the actors connected by links. Meetings consume actors’ resources, but they also contribute toward coordination.

All of the structural elements (e.g., mission tasks, requirements and interdependencies; actor capabilities, skill levels and roles; organization structure, task structure and meeting requirements) of this VDT model are developed by the authors. Such structural elements would clearly be different for each unique organization and process model. VDT also includes several dozen environmental variables with “normal” values determined empirically by prior field

research. These include factors such as the level of uncertainty and noise associated with a project, the inherent propensity of an organization to make errors, and relative concern for performance quality associated with actors at different levels of organizational hierarchy. These and other environmental variables can be changed where appropriate to reflect a wide variety of different organizations and contexts. Other factors can be changed to reflect different organizational designs. For instance, the level of centralization and formalization can be varied by changing design variables. The corresponding VDT model behaviors have been developed empirically.

VDT also includes several performance variables for comparison. In addition to standard simulation measures such as project duration and cost, VDT also includes measures such as levels of rework, coordination and delay, in addition to risk measures keyed to various attributes of importance (e.g., tasks left undone, missed communications, project-level errors). Some of these performance variables are correlated often with one another, whereas others highlight tradeoffs that must be made. For instance, where a project is running behind schedule but on budget, a leader or manager can decide to employ more resources. This often has the effect of increasing the rate of progress but also increasing the rate of expenditure. Other tradeoffs such as those between cost and risk or schedule and coordination require balance in a similar fashion. It is important to note again, the extensive and longitudinal validation of VDT provides considerable confidence that the organizational behaviors emulated by our computational model will reflect well those of operational C2 organizations in the field.

VDT Computational Experiment

Through computational experiments, we emulate the behaviors of a modeled organization as subjected to different conditions (e.g., mission task difficulties, coordination loads, experience levels) and designs (e.g., organizational structures, personnel characteristics, technologies). In general a researcher performing an experiment—whether computational or not—would investigate the background literature and develop a set of hypotheses for testing. In this article we skip the literature review step, for we are interested in illustrating the method of computational experimentation, not the results of a particular experiment per se. Nonetheless, in this section we examine experimentally three factors that should be of interest to the reader: 1) *level of bureaucracy*; 2) *coordination load*; and 3) *knowledge inventory*.

Briefly, the level of bureaucracy pertains to the organizational structure and is operationalized through a combination of VDT constructs (e.g., degree of centralization, level of formalization, lateral information seeking, hierarchical levels, coordination via meetings, team experience). Such constructs for bureaucratic forms are cited widely in the organization studies literature (Scott 2003). Coordination load pertains to the task environment and is operationalized through a different set of VDT constructs (e.g., levels of task interdependency, communication requirements, noise, project-level coordination difficulties). These constructs are grounded similarly in the organization studies literature. Knowledge inventory pertains to the capability of the organization and is operationalized through two VDT constructs (e.g., skill level, capability match). Skill level pertains to how well an actor can perform a certain class of tasks. Capability match pertains to the class or classes of tasks with which the actor has developed skill.

For the experiment, each of these three factors is specified at two levels: high and low. Hence a full-factorial design consists of eight trials, which we designate according to the levels corresponding to the three factors. For instance, the first trial involves High *bureaucracy*, High *coordination load*, and High *knowledge inventory* (HHH). The second trial involves High

bureaucracy, High *coordination load*, and Low *knowledge inventory* (HHL), and so forth. We report on two dependent variables of particular interest in the C2 domain: mission duration and mission risk. As the name implies, mission duration pertains to the elapsed time required for a mission to reach its completion milestone. The importance of speed in warfare is known well, particularly in modern times. Mission risk is measured in VDT as the fraction of assigned mission tasks left incomplete at the end. The completion of mission elements has great bearing on the efficacy of the mission as a whole. Clearly every single task planned for a mission need not be completed for the mission objective to be attained and the mission as a whole to be a success. But the more mission tasks that remain incomplete, the greater the risk to mission effectiveness. Notice that going back to correct deficiencies and complete unfinished mission tasks requires additional time but contributes to efficacy. Hence these two dependent variables *mission duration* and *mission risk* set up a tradeoff between speed and efficacy. The tension between these two performance measures serves to highlight several important tradeoffs between our three experimental factors *bureaucracy*, *coordination*, and *knowledge*.

Experimental Results

Here we report the results of the C2 computational experiment described above. The key results are summarized in Table 1. For each of the eight trials associated with this full-factorial experiment, the table includes measures for mission duration (i.e., measured in days to complete the final milestone) and mission risk (i.e., measured in percentage of mission tasks left incomplete at the end). For instance, the first trial involves High *bureaucracy*, High *coordination load*, and High *knowledge inventory* (HHH). The mission duration for this trial is 347 days, indicating that nearly a calendar year is required to reach mission completion. The mission risk is 36%, indicating that roughly 36% additional resources (e.g., people, time) would be required to complete all mission tasks. The other table entries are reported in the same manner.

Table 1 Experimental Results

Trial	Bureaucracy	Coordination	Knowledge	Duration (Days)	Risk (%)
1 HHH	High	High	High	347***	36*
2 HHL	High	High	Low	361 ***	33***
3 HLH	High	Low	High	335***	37***
4 HLL	High	Low	Low	351***	33 ***
5 LHH	Low	High	High	333***	57 ***
6 LHL	Low	High	Low	347***	55***
7 LLH	Low	Low	High	325 ***	56***
8 LLL	Low	Low	Low	337***	53***

*** significant at the 99% level; ** significant at 95%; * significant at 90%.

The VDT modeling environment uses Monte Carlo techniques to simulate the variation necessary to support statistical analysis. Each performance measure reported in the table represents the mean of 100 simulation trials, and VDT also provides standard deviations for such measures. Hence confidence intervals can be calculated, MANOVA can be conducted, and hypotheses can be tested statistically. For instance, all of the duration results reported in the table are statistically different, at the 99% level, than mission duration of the baseline model (mean =

356, $sd = 6$), as are all of the risk results (mean = 0.35, $sd = 0.059$) except for the first trial (HHH), which is significant at the 90% level. Due to the expository nature of this article, we omit other routine analysis such as testing for main and interaction effects, conducting contrasts and the like. The point here is, such statistical analyses can be conducted through computational experimentation research.

In terms of interpretation, notice the maximum mission duration corresponds to the second trial (HHL). Where bureaucracy is high and coordination load is high but knowledge inventory is low, the mission requires the most time. Hence aspects of bureaucracy contribute toward extended mission execution as do requirements for heavy coordination. The contribution of knowledge to mission execution is intuitively appealing, and we show here how such contribution can be measured computationally. Alternatively, the maximum mission risk corresponds to the fifth trial (LHH). Where bureaucracy is low but coordination requirements are heavy, risk increases. This reveals the tension associated with bureaucracy: although this form of organization contributes to extending mission duration, it reduces mission risk. The knowledge inventory does not contribute much toward mission risk, but the contribution it makes is subtle. Where knowledge inventory is high, work progresses rapidly, for people and organizations know what they are doing. However, such rapid progression of work actually leaves less time to correct mistakes that are made and attend to communications that are missed. The net effect is a very small increase of mission risk associated with high knowledge levels. Anecdotally this bears some resemblance to overconfidence. The minimum mission duration (325 days) corresponds to the seventh trial (LLH), and the minimum mission risk (33%) corresponds to the fourth (HLL). Notice these minimum trials reflect mirror images of their maximum counterparts from above in terms of factor levels (cf. HHL vs. LLH for duration, and LHH vs. HLL for risk). This is appealing intuitively and suggests excellent consistency between the eight experimental trials.

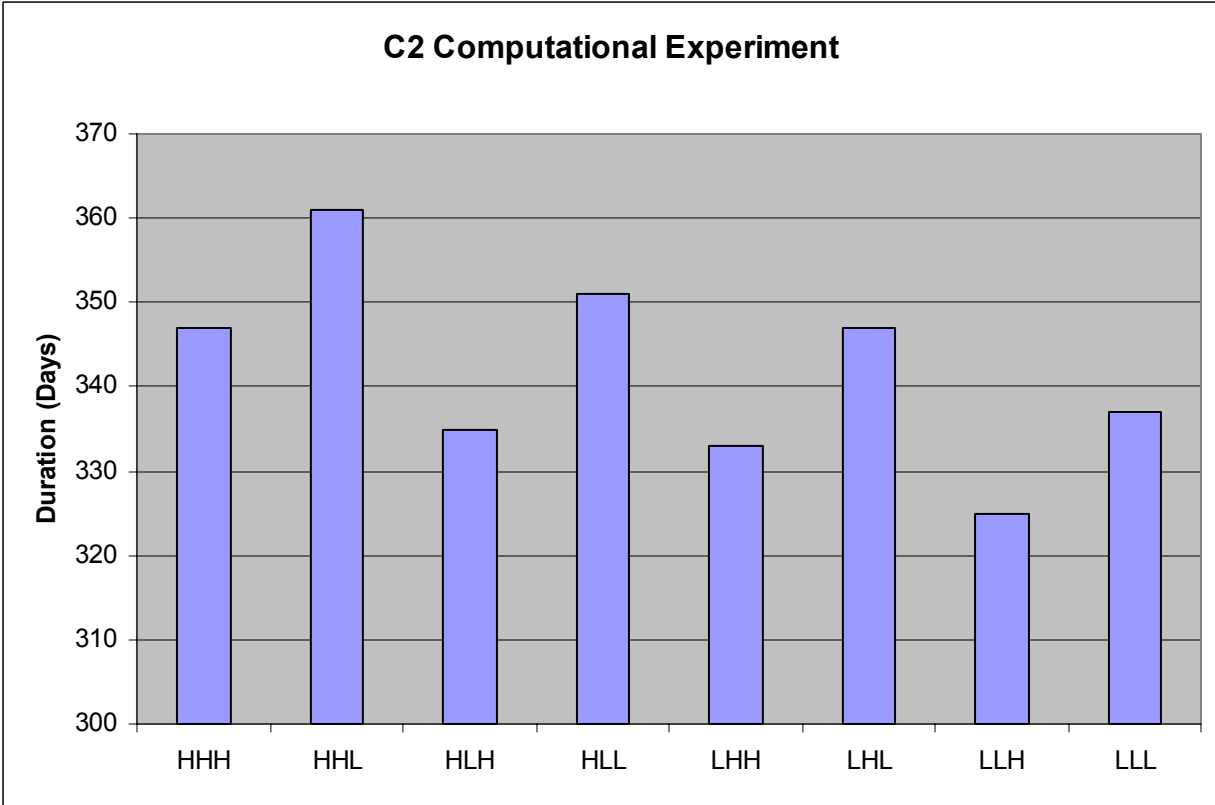


Figure 4 Mission Duration Results

Experimental results are also presented in Figure 4 for additional interpretation. Here we present a bar chart depicting mission duration associated with each trial. As noted above, in each case bureaucracy adds to mission duration, as does coordination load, but knowledge inventory reduces the time required. Comparing visually the bars in this chart serves to reinforce the relationships summarized in the table. Figure 5 presents the complementary bar chart depicting mission risk associated with each trial. Notice the considerable difference attributable to bureaucracy when compared to the other factors examined in this experiment. As above, comparing visually the bars in this chart serves to reinforce the relationships summarized in the table.

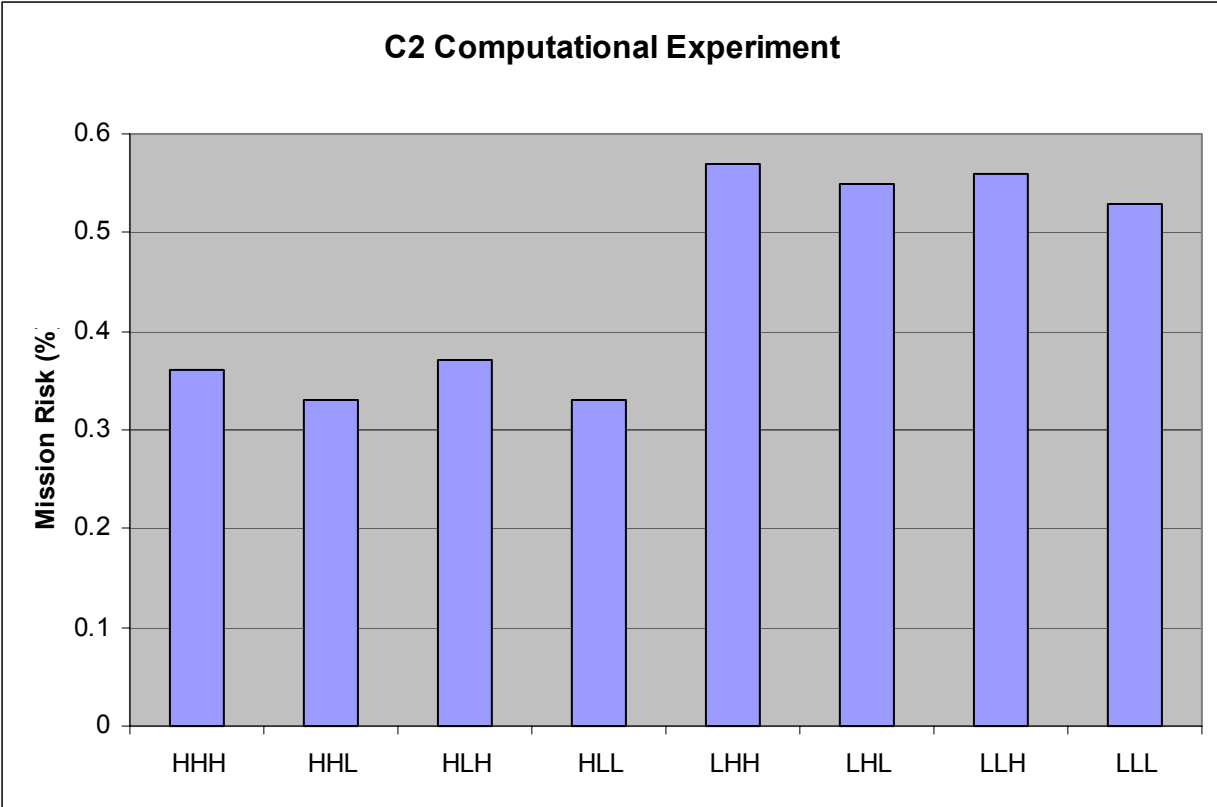


Figure 5 Mission Risk Results

Finally, Figure 6 presents a scatter chart plotting the relative coordinate positions of all eight trials with respect to both mission duration (X axis) and mission risk (Y axis). This view is useful to gauge both the individual and combined effects of the three experimental factors. For instance, the graph delineates clearly the two different levels of mission risk associated with bureaucratic and non-bureaucratic organizational form; all of the former points plot below the 40% line, and all of the latter plot above the 50% risk mark. But the individual contributions of coordination load and knowledge inventory are not quite so clear in this representation. Instead, both coordination load and knowledge inventory interact together to create the dispersion across the duration axis. For instance, where coordination load is high and knowledge inventory is low (i.e., for a given level of bureaucracy), the effect on mission duration is striking.

Interpreting further these figures, where mission speed is of primary concern to a leader, then the bureaucracy represents an inferior form of organization than counterparts with lesser degrees of centralized decision-making, formalized procedures and vertical information flows. Alternatively, where mission risk is primary, then the bureaucracy represents a superior organizational form. This reflects a fundamental tradeoff between performance measures and organizational design, as conceptualized generally in terms of Contingency Theory (Lawrence and Lorsch 1967). Further, high coordination loads place considerable stress on C2 processes. Coordinated defenses and attacks represent commonplace in warfare today, but the associated task interdependencies and communication requirements can actually debilitate C2. Additionally, where knowledge inventory is low, the organization needs particular help to ensure mission duration does not extend greatly and mission risk does not increase appreciably. It is important to

understand the knowledge of an organization and to compare such knowledge with mission demands. But how many commanders today possess the insight, much less the tools, to assess knowledge inventory as such. Clearly primitive measures used today such as *manning* and *readiness* fall hopelessly short.

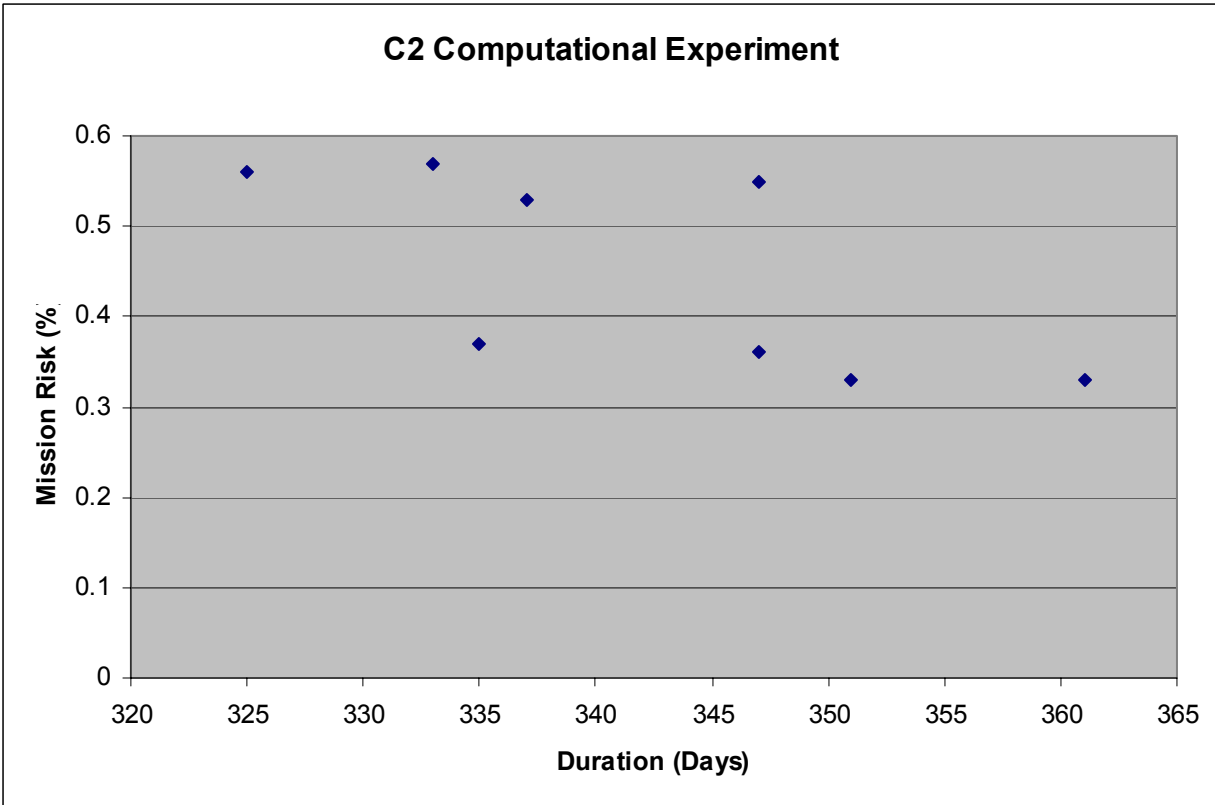


Figure 6 Mission Duration vs. Risk Results

Conclusion

A chasm exists between laboratory and field methods in C2 research. These methods are complementary but used rarely in combination. This expository article describes a research approach that bridges such chasm: computational experimentation. Computational experimentation mitigates the weakness of both laboratory and field research, yet it has its own limitations and appears suited best as a complement and not a replacement. To illustrate the power and potential of computational experimentation, we describe an implemented agent-based modeling environment called VDT. VDT benefits from accumulated research over two decades and extensive external validation. We employ this modeling environment to represent and emulate the behavior of a high-level C2 organization. Using a full-factorial experimental design, we illustrate computational experimentation through controlled manipulation of key factors associated with organizational and technological design (i.e., bureaucracy level, coordination load, knowledge inventory). This illustration includes discussion of rich operationalized constructs used to characterize a diversity of C2 organizations, task environments and

performance measures. The experimental results highlight complex interactions between design factors, and they suggest fundamental tension and decision tradeoffs between important performance measures such as mission duration and risk.

We illustrate through the article how computational experimentation bridges the chasm between laboratory and field methods. The baseline, high-level C2 model discussed above was developed over the course of a week, and the computational experimentation required roughly the same amount of time. Of course, the C2 organization we modeled is fictitious, and our model represents such organization at an admittedly high level. Also, we are very experienced with modeling in general and VDT in particular. But what would be the time, effort and cost to execute in the laboratory a full-factorial experimental design such as described in this article? How would a researcher in the field even begin to establish the kinds of controls and manipulate the kinds of design factors exhibited in this study?

We also illustrate through the article how diverse factors such as organizational form, coordination requirements and organizational knowledge play important roles in terms of C2 performance. Understanding when the bureaucracy is relatively important and how this rigid organizational form can impact negatively mission speed is important for C2 practice today, as is the performance impact of coordination load associated with mission task interdependencies. The critical role played by knowledge inventory manifests itself clearly in our experimental results. Knowledge is key to effective work, and effective work is key to organizational performance. The results suggest leaders and managers should assess their knowledge inventory before embarking on a mission, and they should adjust the manner in which each mission is undertaken (e.g., in terms of organizational form, coordination requirements) on that basis.

The article also leads to natural topics for future research along these lines. Although the high-level C2 organization modeled in this study is representative of such organizations in practice, we do not claim to have experimented—even computationally—with an operational organization. A logical future study would take the VDT modeling environment into the field and model such an operational C2 organization. Once the VDT model has been validated to emulate the key behaviors of the operational organization in the field, then any number of different experimental designs (e.g., full-factorial, Latin Squares, blocking with replication) can be executed computationally to develop results as rich as those presented in this article, but further pertaining to an operational C2 organization in the field.

Another logical future study would employ computational experimentation to narrow down the range of promising factors, conditions and scenarios to test in the laboratory. In other words, computational experimentation can be used to inform laboratory experimentation and hence focus laboratory efforts on those conditions that appear most promising in results of computational experiments. Such combination of computational and laboratory methods could inform further and perhaps guide better A2C2 experimentation, for instance. Likewise, results of computational experiments can be used to guide field research and help investigators focus on factors that show performance leverage through experimentation. Such combination of computational and field methods could inform further and perhaps guide better use of technologies such as the Navy's Battle Force In-port Trainer (BFIT), for instance. Developing feedback mechanisms and routines in C2 to recognize when the organizational and technological design matches a particular mission represents a related topic for continuing this line of research. Analytical and fieldwork to develop and operationalize new constructs such as *knowledge inventory*, which is critical for mission success and offers potential for alerting commanders to

abundances as well as deficiencies in organizational capabilities, follow logically from this line of research as well.

Building upon the VDT constructs introduced in this article, one day researchers may even develop techniques for design optimization (e.g., organizational form and technological artifacts) based on mission objectives (e.g., speed vs. risk) and constraints (e.g., coordination load, knowledge inventory). Leaders, managers and researchers may develop the capability to design organizations, work processes and technologies using computational techniques comparable to those employed for designing airplanes, bridges and computers. That day is not yet here. But through research along these lines, we can both foresee and accelerate it.

References

- Bernard, H.R., *Handbook of Methods in Cultural Anthropology* Walnut Creek, CA: Altamira Press (1998).
- Box, G.E.P., Hunter, W.G. and Hunter, J.S., *Statistics for Experimenters: an Introduction to Design, Data Analysis and Model Building* New York, NY: Wiley (1978).
- Campbell, D.T. and Stanley, J.C. *Experimental and Quasi-Experimental Designs for Research*. Chicago, IL: Rand McNally (1973).
- Carley, K.M. and Prietula, M.J., (Eds.), *Computational Organization Theory* Hillsdale, NJ: Lawrence Erlbaum Associates (1994).
- Carley, K.M., "On Generating Hypothesis Using Computer Simulations," *Systems Engineering* 2:2 (1999).
- Carroll, T. and Burton, R.M., "Organizations and Complexity: Searching for the Edge of Chaos," *Computational & Mathematical Organization Theory* 6:4 (December 2000), pp. 319-337.
- Cheng, C.H.F., and Levitt, R.E., "Contextually changing behavior in medical organizations," *Proceedings of the 2001 Annual Symposium of the American Medical Informatics Association*, Washington, DC (Nov 3-7, 2001).
- Chaing, A.C., *Fundamental Methods of Mathematical Economics* (Third Edition) New York, NY: McGraw-Hill (1984).
- Christiansen, T.R., *Modeling Efficiency and Effectiveness of Coordination in Engineering Design Teams* doctoral dissertation, Department of Civil and Environmental Engineering, Stanford University (1993).
- Cohen, G.P., *The Virtual Design Team: An Object-Oriented Model of Information Sharing in Project Teams* doctoral dissertation, Department of Civil Engineering, Stanford University (1992).
- Cook, T.D. and Campbell, D.T., *Quasi-Experimentation: Design and Analysis Issues for Field Settings* Boston, MA: Houghton Mifflin (1979).
- Denzin, N.K. and Lincoln, Y.S. (Eds.), *Handbook of Qualitative Research* Thousand Oaks, CA: Sage (1994), pp. 500-515.
- Diedrich, F.J., Entin, E.E., Hutchins, S.G., Hocesvar, S.P., Rubineau, B. and MacMillan, J., "When Do Organizations Need to Change (Part I)? Coping with Incongruence," *Proceedings 8th International Command and Control Research and Technology Symposium*, Washington, DC (June 2003).
- Galbraith, J.R., *Organization Design*, Addison-Wesley, Reading, Massachusetts (1977).
- Glaser, B.G. and Strauss, A.L., *The Discovery of Grounded Theory: Strategies for Qualitative Research* New York: Aldine de Gruyter (1967).

Handley, H.A.H., Zaidi, Z.R. and Levis, A.H., "The Use of Simulation Models in Model Driven Experiments," *Systems Engineering* 2:2 (1999).

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

Johnson, R.A. and Wichern, D.W., *Applied Multivariate Statistical Analysis* (Third Edition) Englewood Cliffs, NJ: Prentice Hall (1992).

Kleinman, D.L., Levchuk, G.M., Hutchins, S.G. and Kemple, W.G., "Scenario Design for the Empirical Testing of Organizational Congruence," *Proceedings 8th International Command and Control Research and Technology Symposium*, Washington, DC (June 2003).

Kunz, J.C., Levitt, R.E. and Jin, Y., "The Virtual Design Team: A Computational Simulation Model of Project Organizations," *Communications of the Association for Computing Machinery* 41:11 (1998), pp. 84-92.

Lawrence, P.R. and Lorsch, J.W., *Organization and Environment: Managing Differentiation and Integration* Boston: Division of Research, Harvard Business School (1967).

Lapin, L.L., *Quantitative Methods for Business Decisions with Cases* (Third Edition) San Diego, CA: Harcourt Brace Jovanovich (1985).

Law, A.M. and Kelton, D., *Simulation Modeling and Analysis* (Second Edition) New York, NY: McGraw-Hill (1991).

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

March, J. G., and H. A. Simon, *Organizations*, John Wiley, New York (1958).

Nissen, M.E. and Levitt, R.E., "Agent-Based Modeling of Knowledge Flows: Illustration from the Domain of Information Systems Design," *Proceedings Hawaii International Conference on System Sciences*, Waikoloa, HI (January 2004).

Nogueira, J.C., *A Formal Model for Risk Assessment in Software Projects* doctoral dissertation, Department of Computer Science, Naval Postgraduate School (2000).

Scott, W.R., *Organizations: Rational, Natural, and Open Systems* (Fifth Edition) Upper Saddle River, NJ: Prentice-Hall (2003).

Thomsen, J., *The Virtual Team Alliance (VTA): Modeling the Effects of Goal Incongruency in Semi-Routine, Fast-Paced Project Organizations* doctoral dissertation, Department of Civil and Environmental Engineering, Stanford University (1998).

Thompson, J. D., *Organizations in Action: Social Science Bases in Administrative Theory*, McGraw-Hill, New York (1967).

VDT. The Virtual Design Team Research Group website; URL: <http://www.stanford.edu/group/CIFE/VDT/> (2004).

Wong, S.S. and Burton, R.M., "Virtual Teams: What are their Characteristics, and Impact on Team Performance?," *Computational & Mathematical Organization Theory* 6:4 (December 2000), pp. 339-360.

Yin, R.K. *Case Study Research: Design and Methods* (Second Edition) Thousand Oaks, CA: Sage (1994).