15th International Command & Control Research & Technology Symposium

"The Evolution of C2"

Organizational Structure, Exploration, and Exploitation on the ELICIT Experimental Platform

Topics: 5 or 6

Allan Friedman, Harvard University*****
Ethan Bernstein, Harvard University
David Lazer, Northeastern University

* Point of contact
Allan Friedman
Harvard University
Address: Maxwell Dworkin 110
33 Oxford Street
Cambridge MA 02138
Tel: 617-943-2190,
E-mail: allan_friedman [at] ksgphd.harvard.edu

ABSTRACT
Organizational science has identified the importance of balancing exploration (seeking new information and capacity), and exploitation (effectively using current information and capacity). Similar to the command and control research on organizational agility, this line of research stresses the importance of understanding determinants of performance in an uncertain and complex environment. The structure of the organization and its patterns of communication play an important role in navigating this tradeoff to maximize performance. This paper introduces an extension to the ELICIT (Experimental Laboratory for Investigating Collaboration, Information Sharing and Trust) framework to explore how varying organizational network structure impacts collective performance. We vary the ELICIT framework to capture key aspects of exploration and exploitation in group problem solving. The new ELICIT capabilities are motivated and defined to capture the components of search in a complex environment, organizational network structure and experimental implementation. This research approach utilizes the ELICIT tool to further support and extend research on C2 to understanding the value of edge networks in light of an important area of organizational science.

Keywords: ELICIT, experimentation, exploration, exploitation, C2 organization, collaboration, agility, edge organization

Acknowledgements: The research described in this article was funded by the Command and Control Research Program.

Table of Contents

Acknowledgements

The authors are indebted to Mary Ruddy and the ELICIT team for their incredible assistance and patience in the slow and iterative development process. Thanks also to the staff at Harvard Business School's Computer Lab for Experimental Research and the Kennedy School's Harvard Decision Science Lab.

INTRODUCTION

Through the lens of agility (Alberts & Hayes 2003; 2006), command and control research (C2) has focused on how organizations can cope with the challenges of complexity and uncertainty that characterize the modern institutional environment, both inside and outside the national security context. A subset of the challenge of adapting in a complex environment is the balance between exploitation, or excelling at current tasks, and exploration, which is the ability to adapt to new environments (March, 1991).

This paper explores the importance of balancing exploration and exploitation in the context of C2, and particularly focuses on the challenge of building organizations to optimize this balance. We present an extension to the ELICIT platform that effectively captures this dynamic to test the hypothesis that differences in defined organizational structure can impact how information is gathered and processed by individuals. We discuss the importance of capturing the individual exploration process, as well as how key aspects of the communication network structure can impact the group performance. Finally, we introduce several experimental execution modifications to maximize the power of the experiment.

MOTIVATION: THE ORGANIZATIONAL STRUCTURE OF EXPLORATION AND EXPLOITATION

In thinking about how organizations adapt to complex environments with uncertainty, James March (1991) theorized that organizations must balance exploration and exploitation. Organizations in a changing and challenging environment must devote energy to both learning about their environment and to taking advantage of it. This dynamic has served as a cornerstone in organizational analysis in a wide range of contexts. While definitions vary depending on the context (Gupta, Smith and Shalley, 2006) exploration generally involves attempts to introduce new information, while exploitation leverages existing knowledge for some productive end. Organizations which successfully balance the two have been labeled ambidextrous (e.g., Duncan 1976, Tushman & O'Reilly 1996) and are noted for their superior problem-solving performance, while those that don't risk temporary missteps and/or permanent failure.

A key aspect of this balance between obtaining new information and making use of existing information is the patterns of information flow. Lazer and Friedman (2007) explicitly link structure of communication networks with organizational performance in a problem-solving context. They find that, for complex problems, networks that are more efficient at disseminating information perform better in the short run, but worse in the long run. This is because inefficient networks maintain a diversity of ideas and theories relevant to the problem, and thus are better at exploration of the problem space by supporting a more thorough search of the problem space. Increased communication results in a faster convergence to one solution by driving out all that look inferior, shifting the network to an exploitation bias. However, if the problem environment is

complex, some of these initially inferior approaches may ultimately lead to optimal solution with further exploration.

The question we seek to answer is which communication structure, defined by patterns of interaction between problem solving individuals, will maximize the likelihood of achieving the optimal solution in the shortest possible time. An efficient network is predicted to positively affect information diffusion, which should facilitate the spread of effective strategies, but negatively affect information diversity, which is also positively related to performance. The ELICIT platform allows us to capture the state of awareness of individual problem-solvers as the progress through their task, as well as measure outcomes in terms of individual and collective performance to learn more about this balance.

AGILITY AND AMBIDEXTERITY

The dynamics of exploration and exploitation have obvious relevance to the C2 framing of organizational agility, or the "capability to… meet the challenges of complexity and uncertainty." (Alberts, 2007). A full theoretical mapping from the exploration-exploitation model to the concepts of agility is beyond the scope of this paper, but several brief observations can be made.

In their seminal book on command and control, Alberts and Hayes (2006) argue that agility is actually composed of robustness, resilience, responsiveness, flexibility, innovation, and adaptation. While each of these has been expounded upon elsewhere, it is worth pointing out that robustness and resilience both deal with the ability to maintain effectiveness and focus, which can be mapped to exploitation. Alternatively, one could look at agility as focus and convergence (Alberts, 2007). Focus in an organization is a shared awareness and context, or an "actionable intent." This can be construed as that which is necessary for driving action, or some new component of information. Similarly, if convergence is the movement towards a single purpose and a single entity, one might imbue it with exploitative properties. However, as Alberts makes clear, focus and convergence are best understood properties of disjoint collectives; one can evaluate an organization and measure them, rather than capacities like exploration and exploitation. Moreover, Alberts does not explicitly set them in contrast, unlike the dynamics of exploration and exploitation.

Still, in spite of the lack of an overarching theoretical mapping between the exploration/exploitation dynamic that drives this experimental model and the agility approach to C2, there are common ties. The problem-solving model rests on the idea of self-synchronization, a key aspect of the Edge C2 model. Each actor will receive some information from the world, and have to decide how to explore to seek out more, and when to exploit. The structure of communication will determine how individuals choose to coordinate their efforts, assume responsibility, share key facts that help others and successfully balance the need for exploring and exploiting information. Control of the problem-solving process, like the modern command structure, "is in fact an emergent property". (Alberts, 2007)

BRIEF OVERVIEW OF THE ELICIT MODEL

The United States Department of Defense Command and Control Research Program (CCRP) of the Office of the Assistant Secretary of Defense for Networks and Information Integration (OASD/NII) is engaged in developing and testing principles of organization that enable transformation from traditional hierarchy-based command and control practices toward the transference of power and decision rights to the edge of the organization. The need for agility in Information Age militaries is becoming increasingly important. As discussed in *Understanding Command and Control* (Alberts & Hayes, 2006), in an era of complex, coalition, civil-military operations, understanding how to organize for agility not just within a specific organization but also across differing organizations and cultures is a key to success.

There has been a shortage of formal experimentation data on the efficacy of different C2 organizational approaches. In order to remedy such shortage, the CCRP created and continues to sponsor and maintain the ELICIT experimentation environment. ELICIT is a Java-based software platform that can be used to run multi-user experiments focused on information, cognitive, and social domain phenomena. People participate in experiment sessions mediated by ELICIT by working together in teams that can be configured to reflect different organizational approaches (e.g., Hierarchy, Edge, Caveman, etc.) and that can be subjected to a wide variety of experiment controls and manipulations.

This ELICIT experimentation platform has configurable scenarios that focus on the task of discovering the "Who", "What", "Where", and "When" of a fictitious terrorist plot. Information in the form of "factoids" is provided periodically to each of the participants during an experiment session. The factoids and their distribution are structured so that no one participant receives all the information necessary to perform the task; thus, information sharing is required in order for any participant to be able to determine a solution to the ELICIT problem.

ELICIT provides an instrumented task environment that captures and time stamps participants' information sharing activities. The environment generates detailed transaction logs summarizing such information. These logs, together with participant surveys that can be administered either prior to a trial (for calibration), after a trial, or *in situ*, can be used to measure information sharing, collaboration behaviors and situational awareness, as well as a variety of value metrics including the ability of individuals and teams to correctly identify the future adversary attack and the time required to do so. Considerable research has been conducted to date using ELICIT (Leweling & Nissen, 2007; Powley & Nissen, 2009), and the interested reader is directed to the corresponding references for details and results. The ELICIT experiment platform and processes are significantly enhanced as part of this effort to capture the dynamics of exploration and exploitation in a networked environment. These changes are detailed below.

SEARCH IN A COMPLEX ENVIRONMENT

As mentioned above, a key aspect of this research project is the importance of modeling exploration. Exploration assumes several components, each of which are carefully implemented in the ELICIT software platform. The critical aspects are to understand the problem as a complex problem necessitating myopic search, and a search process that requires subjects to develop and then test specific hypotheses. Together, they create a problem context where individuals have to assess the information available, construct theories and then act on those theories.

The first approach is to recreate the idea of myopic search. (Levinthal and March, 1993). In myopic search, the ability to easily recognize the correct solution is curtailed to reflect cognitive, observational and operational limitations. That is, the search for the solution is bounded to solutions similar to the problem-solver's current perspective. Of course, any nontrivial search through a problem space is to some extent bounded; otherwise, finding the solution would be trivial. Myopia makes sense as a search process, as potential solutions similar to those an individual already understands are more evident and visible, simpler to implement, and easier to assess.

We implement myopic search in ELICIT by altering the factoid sets in two critical ways: the use of "red herrings" and the elimination of "silver bullets." In the original factoid sets, some "noise" factoids were obviously and immediately identifiable as noise. In the revised set, we made the noise factoids more misleading. Now most factoids, if examined alone, might appear relevant to a participant. Moreover, several factoids together might create a plausible theory, diverting a participant from pursuing the correct answer. Should a participant not readily identify a new factoid as noise, they might spend time and effort searching for more information to validate or disprove the theory, discuss this with others, consuming their attention and bandwidth as well, or even register the new theory as their best guess. These red herrings serve as local optima that distract from the search for the global optimum.

With respect to the role of an individual factoid in determining the solution, in previous factoid sets, there was a "silver bullet" factoid which answered a component of the problem—e.g., "The attack will occur at 3pm." Given our search functionality, we removed all silver bullet factoids, so that every solution component requires at least 2, and on average 3, factoids to determine. This requires the use of search and information sharing to solve the problem.

Search alters the paradigm of information distribution. Rather than only receiving factoids at defined times, ELICIT is enhanced so that subjects are able to search for factoids based on single-word queries based on theories they develop, individually and collaboratively, over time. Search is only permitted once per two minute interval. However, search is powerful: search only returns unique factoids (i.e., factoids not already in the subjects" inbox). If search provides no result, subjects are permitted to search again—null results are costless.

Given our search functionality, again for the purpose of verisimilitude, we modified the approach to factoid distribution in two key ways. First, not all factoids are distributed—some are only accessible through search. In a real-world situation, it would be unusual if 12 members of a problem-solving team passively received all the information they needed, over time, to solve a problem. Instead, they would likely have to search for some key factoids that would otherwise not come to them—and this change mirrors that reality. Second, all factoids that are distributed are distributed at the very beginning. This benefits both the statistical power of our runs (as there are no exogenous impacts on the experiment after it is started) and the verisimilitude (as, in real life, there is far less serendipitous receipt of answers).

Enhancing the Experimental Power of Search

In designing this extension, it was important to make sure we were effectively capturing user search decisions in the exploration/exploitation tradeoff. We introduced three other major modifications to the experimental process to further isolate the search process as an observable phenomenon.

First, in a substantial deviation from the original model, we revised the factoid set to create independent sub-problems in the Who, What, Where and When of the problem. In previous ELICIT experiments, a "problem" and its "solution" were defined as the Who, What, Where, ***and*** When of a factoid set. By instead defining each individual component (Who, What, Where, ***or*** When) as a problem, we increase the power of our data by 4 times (i.e., each subject solves 4 x 3 = 12 problems in the course of an experiment session, rather than just 3). Thus, each component of the factoid set problem is made independent of the others—e.g., a factoid for Who does not aid in the solution to What.

Second, subjects were explicitly told that their compensation for experimental participation was directly tied to their performance in solving the problem. Only correct answers would lead to reward, and the longer they had the correct answer—that is, the sooner they entered it into the system--the higher the reward. Subjects were rewarded for each sub-component that was correct, on a minute basis, rather than on the correctness of the entire problem.

Finally, to further prompt the subjects to enter their best guesses frequently, subjects were explicitly encouraged to repeatedly enter their best guesses, and told that there was no penalty for being wrong. To give an additional push, we force progress checks every 5 minutes—the "boss" checking in with the team—although the answers are pre-populated with the subject's answers from previous inputs to avoid frustrating the subjects or wasting time.

These enhancements to the ELICIT platform allow us to have a very accurate and granular picture of the "internal state" of each participant as they navigated the search process. We can combine this with the observed behavior logged from the ELICIT system to get a clear picture of how individuals and the collective group balanced seeking new data and understanding their current information to solve the problem.

ORGANIZATIONAL NETWORKS

Changes to Communication Channels

The research approach is to explore how the organizational structure that defines the patterns of interaction affects individual and group performance. A key driving assumption is that individuals should be as interchangeable as possible to examine the role of the organizational network. That drives several key differences from the original ELICIT approach. First, there are no specific task assignments. Everyone is tasked with solving the Who, What, Where and When components, much like the original Edge network. Similarly, approaches to decision-making and role are left to the discretion of each individual in the experiment. No one is put in charge of anyone, even in the hierarchy described below.

Another distinction is that factoid sharing is exclusively defined by the communication networks: there are no web sites or other "pull" sources of socially archived and sharable information. As with the original ELICIT model, each individual decides when to share a factoid with anyone else, and to which of their potential communication partners they wish to send it.

Two additional means to share information through the patterns of connectivity are added to the experiment platform. First, individuals may annotate factoids they send to one another. We wanted subjects to be able to share theories, not just factoids, and thus enabled free-text annotations of shared factoids. It is important to note that this is not the first ELICIT modification to allow text sharing (Ng, 2008). Not all subjects choose to annotate every factoid. In preliminary experiments, the additional information conveyed ranges from the persuasive, ("This proves it was not alphaland") to the coordinative ("This means we should each search for months in the winter") to the relatively content free ("I think this is important.")

The second new channel of information is the ability to see other subjects' current guess. Since subjects are incentivized to guess frequently (see below), we can think of their current logged guess as their best operating theory of the correct answer. Making this visible can be thought of as an environment where theories about the solution can be shared.

Changes to the Network Structure

The original ELICIT models use two main network structures to represent organizational information flow. In the hierarchy, a single commander coordinates four team leaders, each of whom works with a team of four others. In an Edge network, each individual can interact with any other through the information channels available. Earlier results (Leweling, 2007, Ruddy 2007) suggest that the Edge network can outperform the Hierarchy. The network structures described below seek to further explore the exact structural mechanics of how information flows in the problem-solving environment, and how that information flow might relate to performance.

A primary goal in the network design is to minimize the differences in individual subject experience, varying only the global structure. As a result, we tried to keep the degree distribution, or number of ties for each participant, constant across the networks. That is, each participant should have a similar subjective experience in terms of number of communication partners, both inside and across treatments. This allows us to rule out any story involving information overload or information starvation, at least from first principles. Any observed imbalance in information distribution at the subject level will thus be an emergent phenomenon resulting from the global properties of the network.
We present four main network archetypes that can be explored with the modified ELICIT tool. The first three are presented below in Figure 1, with their statistical properties described in Table 1. Each contain similar local structures, but the global pattern differs. We use 16 nodes rather than 17 nodes for ease of symmetry, with the assumption that one less problem solver will have minimal affect.

a. The first network is built on the '*caveman*' style (Watts, 1999) where small clusters of individuals are connected as loosely as possible. A principle characteristic of the caveman network is local clustering with long average path length between any two actors. This might replicate small teams with intense collaboration bound together in a larger organization, where points of contact rest with fewer actors; most people have common ties.
b. The second network is a modified caveman that is rewired by creating four "shortcut" links across the network, with four fewer local ties. The *rewired caveman* still has a decent amount of local clustering, but a much shorter average path length. Note that every individual is now connected to someone who does not share any common ties.
c. The final network is a more conventional *hierarchy*, but we tried to maintain the degree distribution of the caveman networks as closely as possible. It is important to recognize that, unlike the original ELICIT hierarchy, information can flow around the top of the network through the "captains" without requiring a central coordinator.

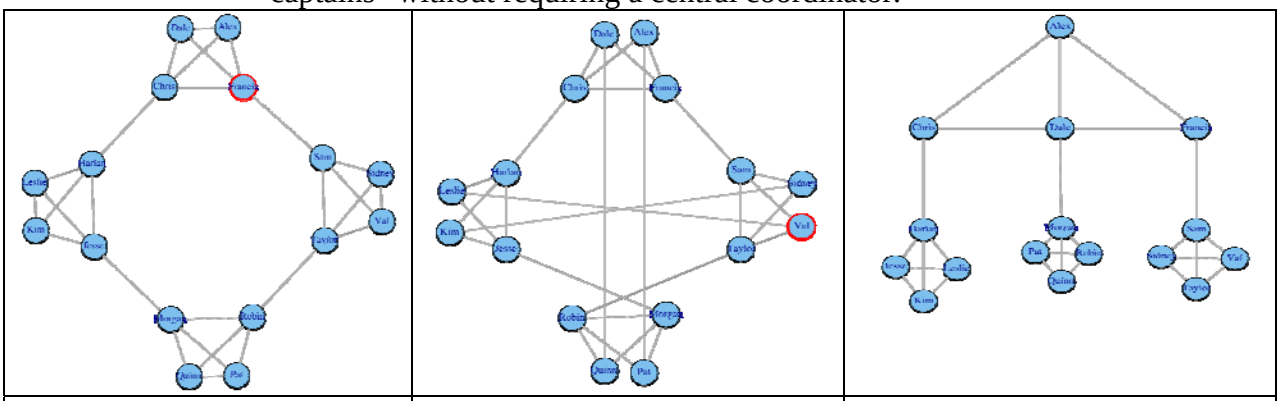

Caveman Network	Rewired Caveman Network	Hierarchical Network

Figure 1 - Three organizational structures with similar local characteristics but with different average path lengths.

	Degree Distribution			Shortest Path Distance Distribution				
Network	**3 ties**	**4 ties**		**1 hop**	**2 hop**	**3 hop**	**4 hop**	**5 hop**
Caveman	50%	50%		21.9%	18.7%	31.2%	15.6%	6.2%
Rewired Caveman	50%	50%		21.9%	34.3%	37.7%		
Hierarchy	62.5%	37.5%		21.1%	14.1%	23.4%	14.1%	21.1%

Table 1 - Statistical properties of the three major organizational structures.

These network structures allow us to compare the effect of a longer average path length, to explore how theories and facts are shared. We can test the findings of Lazer and Friedman (2007) and the assertion that less efficient connectivity can produce higher outcomes by preserving local diversity of ideas, even controlling for the overall number of network ties. The hierarchy offers another perspective, where a full cluster structurally separates each branching team.

It is possible that a 16-node network is not large enough to effectively capture the power of the smaller worlds. The average path length of the "small world" cave is 2.46, compared to the rewired path length of 2.03. Thus, in addition to the three networks described above, we include an environment with the maximum possible average path length. To this end, we also run ELICIT on a 2-dimensional ring lattice, where each actor can only see two others. This creates the longest average path length in any network that is symmetric across individuals and fully connected.

Additional Changes to Experiment Protocols

The experiment design calls for making a very large number of experiment runs (65 experiment runs by the end of April, 2010). Thus, the experiment protocol is further adjusted to ensure that multiple experiment runs can be conducted in a single subject session. The subject pool at Harvard can be scheduled for two hour blocks, so the experiment processes are streamlined until three rounds can be conducted with each group of subjects. The streamlined subject protocol is as follows:

1) A pre-briefing video is used to insure that subject instructions are consistent across rounds.
2) A pretest (first round) is conducted in which each subject receives all the information needed to identify the Who, What, Where and When of a fictitious terrorist scenario. This is used as a baseline for each subject's ability to perform this type of task independent of the organizational structure
3) A second round is conducted that is an actual experiment cell.
4) A third round is conducted that is an actual experiment cell. (Note that different factoid sets are used in each round.)
5) A survey is conducted to gather demographic information about each subject and to obtain additional information about participant's situational awareness. (Note that subjects consistently indicated that they enjoyed participating in the experiment.)

6) Consistency of the experiment protocol is achieved by using a multi-page experiment moderator checklist.

To comply with Harvard Institutional Review Board (IRB) requirements for subject payment, to ensure subject motivation and to ensure that subjects communicate their situational awareness to the experiment platform frequently, subjects receive incentive payments. The incentive payments are based on the cumulative number of minutes that a subject has the correct answer to each part of the Who, What, Where and When of the fictitious terrorist scenario. Subjects register their current best answer in ELICIT using a new variation of the Identify action. Since the when portion of the answer consists of 4 subparts (month, day, hour, AM/PM), there are a total of 7 pieces of information (who, what, where, month, day, hour, AM/PM) on which subjects are measured.
To streamline payment logistics, ELICIT is enhanced to automatically determine if a subject currently has the correct answer as his/her best guess for each of the 7 components and to automatically keep a running total of the accrued payment for each subject.

CONCLUSION

Development of and experimentation with ELICIT is an ongoing activity of the CCRP. A recent CCRP-sponsored effort resulted in the development of new ELICIT platform capabilities to enable experiments in key aspects of exploration and exploitation in group problem solving. The modifications to the platform allow an experimental design to capture the process of search in different network organizations while still reflecting the original designs of the ELICIT program. Under this revised research framework, we have been able to run 65 experimental trials of C2-relevant task simulation with over 380 individuals to better understand the cognitive and social impacts of C2 approaches and organizational structure.
As this paper went to press for the 15th annual ICCRTS, the data collection phase of the project has just been completed. The authors will make their final results, conclusions and data available on the CCRP website http://www.dodccrp.org/html4/elicit.html.
A final note on the experimental process in the ELICIT platform. We have received overwhelmingly positive feedback from our subjects on the ELICIT tool. They find the game fun, the social interaction component engaging, and the operational continuity provided by the online platform a hassle-free experience.

REFERENCES

Alberts, D. S., & Hayes, R. E. (2003). *Power to the edge : Command and control in the information age*. Washington, DC: Command and Control Research Program.

Alberts, D. S., & Hayes, R. E. (2006). *Understanding command and control*. Washington, D.C.: CCRP Publications.

Alberts, D.S. (2007) Agility, Focus and Convergence: The Future of Command and Control. *The International C2 Journal. 1:1 1-30.*

Ng, E.C., Thunholm, P., Cheah, M., Tan, K.Y., Chua, N. (2008), Exploring Alternative Edge versus Hierarchy C2 Organizations Using the ELICIT Platform with Configurable Chat System, *Proceedings International Command & Control Research & Technology Symposium, Seattle, WA.*

Duncan, R. B. (1976). The ambidextrous organization: Designing dual structures for innovation. In R. H. Kilmann, L.R. Pondy and D. Slevin (eds.), *The management of organization design: Strategies and implementation.* New York: North Holland: 167-188.

Gupta, A. , Smith, K., and Shalley, C.E., (2006) "The interplay between exploration and exploitation." Academy of Management Journal, 49: 693–706.

Lazer, D., Friedman, A. (2007) The Network Structure of Exploration and Exploitation. *Administrative Science Quarterly*. 52:667-694.

Leweling, T. A., & Nissen, M. E. (2007). Hypothesis testing of edge organizations: Laboratory experimentation using the ELICIT multiplayer intelligence game. *Proceedings International Command & Control Research & Technology Symposium,* Newport, RI.

March, J., (1991) "Exploration and exploitation in organizational learning." *Organization Science*, 2: 71–87.

Powley, E. H., & Nissen, M. E. (2009). Trust-mistrust as a design contingency: Laboratory experimentation in a counterterrorism context. Washington, DC.

Ruddy, M., (2007). ELICIT – The Experimental Laboratory for Investigation Collaboration, Information-Sharing and Trust, 2007, *Proceedings of International Command & Control Research & Technology Symposium,* Newport, RI.

Levinthal, D.A., & March, J.G. (1991) The Myopia of Learning. *Strategic Management Journal.* 14:1 95-112.

Tushman, M. L. and O'Reilly, C. A. (1996). The ambidextrous organization: managing evolutionary and revolutionary change. *California Management Review*, 38: 1- 23.

Watts, D. (1999) "Networks, dynamics, and the small-world phenomenon." American Journal of Sociology, 105: 493–527.