Destabilizing Dynamic Covert Networks

Kathleen M. Carley
Matthew Dombroski
Max Tsvetovat
Jeffrey Reminga
Natasha Kamneva
Carnegie Mellon University

Contact:
Prof. Kathleen M. Carley
Institute for Software Research International
Carnegie Mellon University
Pittsburgh, PA 15213

Tel: 1-412-268-6016
Fax: 1-412-268-1744
Email: kathleen.carley@cmu.edu

Modeling and Simulation

This paper is part of the Dynamics Networks project in CASOS at CMU. This work was supported in part by the Office of Naval Research (ONR), United States Navy Grant No. 9620.1.1140071 on Dynamic Network Analysis under the direction of Rebecca Goolsby and Grant No. 1681.12.1140053 on Adaptive Architecture under the direction of Bill Vaughn. Additional support on measures was provided under the Darpa project on INSIGHT (Stands for: Interpreting Network Structures to obtain Intelligence on Groups of Hidden Terrorists) DAAH01-03-C-R111.

Citation: Kathleen M. Carley, et al. 2003, "Destabilizing Dynamic Covert Networks" In *Proceedings of the 8th International Command and Control Research and Technology Symposium*. Conference held at the National Defense War College, Washington DC. Evidence Based Research, Vienna, VA.

Destabilizing Dynamic Covert Networks

Abstract:

Most commanders, politicians and intelligence agents have at least an intuitive understanding of hierarchies and how to affect their behavior. However, covert organizations, such as terrorist organizations, have network structures that are distinct from those in typical hierarchical organizations. In particular, they tend to be more cellular and distributed. This makes it difficult to apply the lessons of experience in determining how best to destabilize these groups. This problem is further compounded by the vast quantities of, yet incomplete, information. What is needed is a set of tools and an approach to assessing destabilization strategies that takes these difficulties in to account and provides analysts with guidance in assessing destabilization tactics. Such an approach is forwarded in this paper. In addition, initial lessons learned are discussed. The particular approach is extensible and scales well to groups composed of 1000's of members.

Support: This research has been supported, in part, by the National Science Foundation IGERT in CASOS, the Office of Naval Research (ONR), United States Navy Grant No. 9620.1.1140071 on Dynamic Network Analysis and Grant No. 1681.12.1140053 on Adaptive Architecture, DARPA, the NSF under the IGERT award in CASOS - NSF IGERT 9972762, and the center for Computational Analysis of Social and Organizational Systems. The views and results expressed herein are solely the responsibility of the authors and do not represent the official views of the Office of Naval Research, DARPA or the National Science Foundation.

Approaches to Assessing Destabilization Tactics for Dynamic Networks

Most commanders, politicians and intelligence agents have at least an intuitive understanding of hierarchies and how to affect their behavior. However, covert organizations, such as terrorist organizations, have network structures that are distinct from those in typical hierarchical organizations. A key feature of covert networks is that they are cellular and distributed. Consequently, the lessons of experience held by these decision makers may not be applicable. Reasoning about how to attack dynamic networked organizations (Ronfelt and Arquilla, 2001), let alone figuring out how they are likely to evolve, change, and adapt is terribly difficult. What is needed is a series of tools, techniques, and models for collecting data on and reasoning about these covert networks even in the face of overwhelmingly incomplete information.

To understand the dynamics of covert networks, and indeed any, network we need to understand the basic processes by which networks evolve. Moreover, we have to evaluate destabilization and surveillance strategies in the face of an evolving network and in the face of missing information. To ignore either the dynamics or the lack of information is liable to lead to erroneous, and possibly devastatingly wrong, policies. Taking in to account both the dynamics and the lack of information should engender a more informed approach to answering various policy questions. Key questions might include "what is the size and shape of the covert network", "how does the nation in which the covert network exists impact its form and ability," and "if we do x to the covert network what is likely to happen?"

Two approaches that could be applied to the study of covert networks are traditional social network analysis and multi-agent modeling (particularly a-life). However, both of these approaches are severely limited. Traditional SNA is limited in that it only considers the linkage among people, is concerned with non-adaptive systems, and most measures have been tested only for small (< 300 node) networks. Multi-agent modeling uses very simple unrealistic agents who, although they adapt, move about only on a grid and don't take actual networks in to account. This paper proposes the use of a third approach – dynamic network analysis.

Approaches to Assessment

Dynamic Network Analysis (DNA) extends the power of thinking about networks to the realm of large scale, dynamic systems with multiple co-evolving networks under conditions of information uncertainty with cognitively realistic agents (Carley, 2002b). DNA has been made possible due to three key advances: 1) the meta-matrix (Carley, 2002a; Krackhardt and Carley, 1008) connecting various entities such as agents, knowledge and events, 2) treating ties as "variable" and so having a weight and/or probability, and 3) combining social networks with cognitive science and multi-agent systems to endow the agents with the ability to adapt (Carley, 2002c). In a meta-matrix perspective a set of networks connecting various entities such as people, groups, knowledge, resources, events, or tasks are combined to describe and predict system behavior. In variable tie perspective, connections between entities are seen as ranging in their likelihood, strength, and direction rather than as being simple binary connections indicating exclusively whether or not there is a connection. Finally, the utilization of multi-agent network models enables the user to reason about the dynamics of complex adaptive systems. In particular, these computational models combine our understanding of human cognition, biology, knowledge management, artificial intelligence, organization theory and geographical factors into a comprehensive system for reasoning about the complexities of social behavior.

A key feature underlying this work is a dynamic approach to the co-evolution of agents, knowledge, tasks, organizations and the set of inter-linked networks that connect these entities. Multi-agent network modeling is used to capture the complexities by which who people know influences what they know and so what they can do and what organizations they join. Changes at each unit of analysis, person to group to organization to society impact changes at the next; however, the rate of change decreases and the size of the change's impact increases as unit size increases. Another feature is that each agent (and indeed each unit) has transactive knowledge – knowledge of who knows who, what, is doing what, and is a member of what. This knowledge is typically incomplete, sparse, and potentially wrong. However, the actions of the agents are based on their perception of the network not the actual network. Cognitive, social, task, and cultural constraints limit what entities are present, what/who can be connected to what/who, when and how those connections can change, when new entities (such as new agents) can be added or old one's dropped, and so on.

Proposed Approach

The basic approach that we use to assess destabilization tactics is the following:

1. Identify key entities and the connections among them.
2. Identify key processes by which entities or connections are added or dropped, or in the case of connections, changed in their strength.
3. Collect data on the system (covert network).
4. Determine performance characteristic of existing system.
5. Determine performance characteristics of possible optimal system.
6. Locate vulnerabilities and select destabilization strategies.

7. Determine performance characteristics in the short and long term after a destabilization strategy has been applied.

Some comments on this approach are warranted. First, the result of this process is an evaluation of both system vulnerabilities and the impact of attacking those spots, with some estimate of the robustness of the results in the case of missing information. By providing both the vulnerabilities and the impact of attack, the analyst can use this information to consider the possible ability of these attacks to effect other outcomes other than the specific performance characteristics examined. Second, the process as described above is very general. We have instantiated at this point, and will describe, a relatively simple form of this process. It is important to note that the approach is broader than this simple instantiation. It is in this sense that we say that the approach is extensible.

Application of Approach

We now describe the instantiation of each step in this process. In order to make this illustration more concrete, as we move through each step draw from data collected on an embassy bombing in Tunisia. We refer to this as the embassy bombing data set (EB data set). This data was collected from open source files, such as newspaper reports, by Connie Fournelle at Alphatec.

Identify key entities and the connections among them

For covert networks a number of entities appear to be key: people (agents), knowledge, resources, events, tasks, groups, and countries. This is a complex set of entities. In many cases, detailed data is not known on all such entities. For the sake of exposition, and without loss of generality, in this paper we utilize a smaller set of entities: people, resources, and tasks. We note that in making this simplification we are assuming a one-to-one mapping of knowledge and resources, and another for events and tasks. Thus, if we have data on a person having or using a resource we assume they have the knowledge to employ it. Similarly, if we have data on a person engaging in an event we assume they are engaged in the implicit task and if doing the task are present at the associated event. Thus, e.g., blowing up and embassy (task) and being present when the bomb is set to blow up the embassy are treated as the same entity. This gives the meta-matrix defined in table 1. As a practical note, it is important to recognize that when collecting real data it may be necessary to list some of people by the role they fill, e.g., pilot, or some characteristic, unidentified red haired male, rather than their name or id. Similar types of generalizations may be needed for the other entities.

Table 1. Meta-matrix for covert networks

	People	Resources	Tasks
People	Social network	Capabilities network	Assignment network
Resources		Substitution network	Needs network
Tasks			Precedence network

In addition we specify a number of relations among the entities. Note, there may be many relations among entities; e.g., people can both provide money and provide information to each other. For the sake of exposition, and without loss of generality, in this paper we utilize a reduced set of relations as described in table 1. In this case, we are combining all types of connection in the same cell of the meta-matrix and treating it as a single relation that can vary in strength and direction. As a practical note, different connections will come from different sources and may have different levels of accuracy or strength. At the moment, the system does not distinguish the strength or frequency of the tie from the trust in the source. Thus, the meta-matrix just represents a composite of information from all sources.

Identify key processes by which entities or connections are added or dropped, or in the case of connections, changed in their strength.

Logically, the system can change be adding or dropping nodes or relations. A node can be, given table 1, a person, resource, or task. A relation can be the connection between two nodes. Further, unlike nodes, we can talk of change in the strength of a connection. A number of key processes in covert networks affect these types of changes. Key processes affecting node change include: recruitment; the removal (death, isolation, etc.) of a person; change in mission (and so the addition or deletion of tasks); change in technology (and so the addition or deletion of tasks and resources); the consumption of resources; and the purchase/creation of resources. Key processes affecting the change in relations are re-assignment of personnel, training, co-work assignments, and evolution of friendships/communication structure.

Since we are focused on the destabilization of covert networks particularly in the very short term; e.g., two years. As a first step, we make the simplifying assumption that the rate of change in the mission and technology is much slower than the other changes. Under such an assumption we can treat the mission and technology as fixed and ask, how to destabilize a covert network with a specific mission and specific available technology.

Before continuing it is necessary to specify how the extant mission and technology are captured in the meta-matrix representation. In this case, the mission is characterized by the precedence and needs networks. The technology is characterized by the substitution and needs networks. By limiting this study to the case of a fixed mission and technology, we are in effect saying that the number of resources and tasks do not change. Moreover, we are assuming that the connections in these three networks (substitution, needs and precedence) do not change.

Given these assumptions, the only entities that can be added or dropped are agents. The connections that can change in strength are those between two agents, agents and resources, agents and tasks. Previous work indicates that changes in

Collect data on the system.

Data can be collected from various sources. For this paper, we utilized data in publicly accessible archives, newspaper reports, and professional journals. Although we have several data sets, for illustrative purposes we will use a reduced form of the embassy bombing data (EB data set). Some characteristics of this data are described in table 2. This data set is very small, but it does contain all three entities described before. In creating this reduced form the following simplifying assumptions were made. First, regardless of how the people were connected we simply identified that they had a connection. Second, we combined resources and skills – treating them all as resources. Finally, although the system can handle non-binary relations we coded only the presence or absence of a relation. To illustrate these networks, only the social network is shown in Figure 1.

Table 2. EB data set characteristics

	People	Resources	Tasks
People			
Number of nodes	16x16	16x8	16x5
Resources			
Number of nodes		8x8	8x5
Tasks			
Number of nodes			5x5

Figure 1: EB Social Network

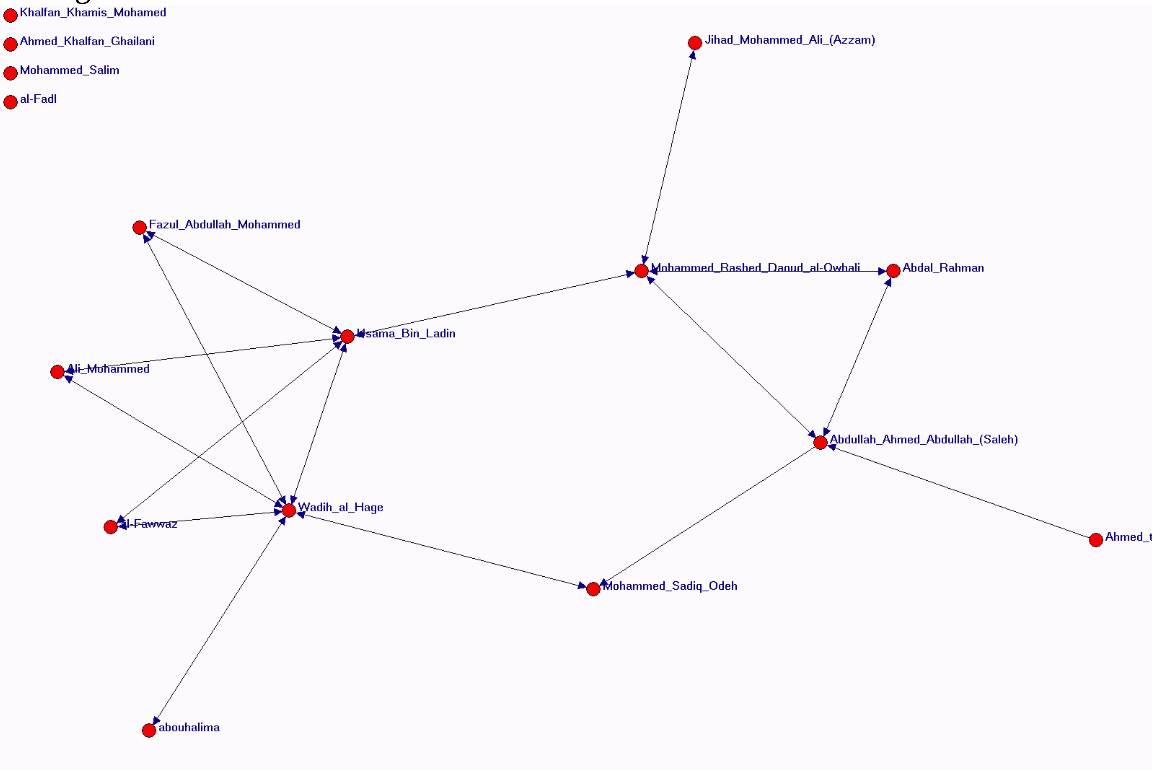

Determine performance characteristic of existing system.

For covert networks, there are many possible indicators of performance that make sense: ability to invoke terror, inter-arrival rate of attacks, severity of attacks, ability to spread information, and so on. For many of these, it is difficult to get sufficient information to validate the model. Further, for some indicators, such as the ability to invoke terror, there are contributing factors that require an assessment of the environment or the object of the attack. This makes the determination of performance potentially intractable.

To address this problem we borrow on the work in organization science. Note, for organizations and teams in general it is difficult to assess performance directly. Hence, analysts have used a variety of indicators of performance including self-assessment, efficiency or effectiveness as perceived by a subject matter expert, financial state, longevity, time to complete a task, or indicators on hypothetical or stylized tasks. In this paper, we use indicators on stylized tasks. In particular, we use two measures in this paper: accuracy of performance on a stylized task and rate of information diffusion. We selected these performance measures as they have previously been shown to reflect team behavior and to map on to actual average performance of teams (Lin & Carley, 2003; Carley & Krackhardt, 1999; Carley, 2002c). Performance as accuracy is measured as average accuracy of the system in solving a suite of binary classification tasks. Performance as diffusion is measured as the average time it takes all members of the system to "learn" a new piece of information.

To measure performance, we take the extant system and simulate it using DyNet. DyNet is a multi-agent network system for assessing destabilization strategies on dynamic networks. DyNet uses the Construct code for assessing information diffusion and accuracy (see for details Carley, 1991, 1999; see also for description of the binary classification task, Carley & Svoboda, 1996). In simulating the system, a knowledge network for the system is given to DyNet as input. We define knowledge here as the individual's knowledge about who they know, what resource they have, and what task they are doing. We make the simplifying assumption that each agent knows about the complete set of available persons, resources and tasks and has no knowledge of what others know. Due to the level of granularity of the data, the alternative assumption that each agent has perfect knowledge of who knows whom, who has what resources, and who is doing what tasks, has little impact on the results.

DyNet is intended to be a desktop system that can be placed in the hands of intelligence personnel, researchers, or military strategists. Through hands-on what if analysis the analysts will be able to reason in a what –if fashion about how to build stable adaptive networks with high performance and how to destabilize networks. Using DyNet we have simulated a variety of covert networks, including the EB network.

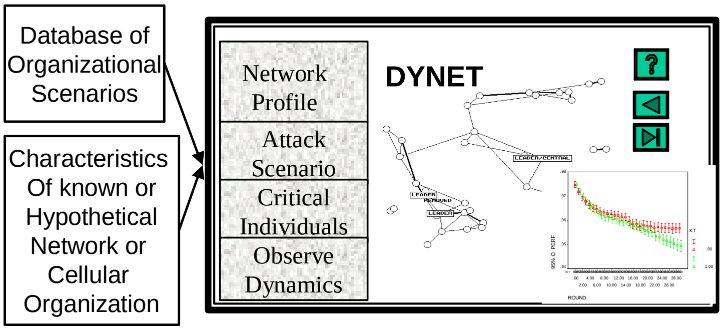

Figure 2. DYNET: A desktop tool for reasoning about dynamic networked and cellular organizations

We ran sensitivity tests examining how differences in the structure of the covert network would impact the overall ability of the network to respond and to be destabilized. The results suggest that isolation of entire cells is more devastating than the isolation of specific individuals. Further, when individuals are isolated, covert networks "heal" themselves. New leaders emerge. Moreover, effective destabilization strategies do not require having full information about the network. Rather than detail these results, we focus on the entire process of analysis.

Determine performance characteristics of possible optimal system.

By identifying the mission and technology constrained portions as relatively fixed components of the extant system, at least in the short run, we open the possibility to locating the optimal form or structure of the rest of the system. We define the organizational design as the set of cells in the meta-matrix that can be varied in the short run – the social networks, the capabilities network, and the assignment network. The system is optimized if the ties in this network are arranged such that they minimize vulnerabilities. We define a system to have the optimal organizational configuration or design if vulnerabilities due to one or more of the following are minimized: distribution of resources, distribution of communication ties, and workload. Previous work indicated that high performance and adaptive systems tended to exhibit a high level of congruence, or match, between what resources were needed for a task and the availability of those resources (resource congruence) and in who needed to communicate in order to do the task and who actually communicated (Carley, 2002c). Further, organizations typically exhibit better performance and have fewer problems with personnel if workload is evenly distributed. Using heuristic based optimization tools, such as simulated annealing, we locate the organizational design that optimizes one or more of these criteria.

There are two ways in which the optimized code can be used. First, if the current design is far from optimal it may not be worth destabilizing at all. Since destabilization involves the removal of critical nodes, the comparison of the relative difference from the optimum of the "destabilized" organization and original provides an indicator of the potential relative impact of the destabilization.

In this paper, we are optimizing purely on the basis of resource congruence. This is sufficient for the sake of illustration. Resource congruence is defined as:

Compute the network's need for resources in order for agents to be able to complete tasks:

```
   let Need = (AR'*AT) - RT
   let d = card{ (i,j) | Need(i,j) < 0 }
   let d = d / (|R|*|T|), which normalizes d to be in [0,1]
```

Then Resource Congruence = 1 – d, s.t., a 1 means complete congruence and a 0 complete incongruence. For the organization examined, the original resource congruence was .475. For the optimized structure it was .8. Note, in doing this optimization the needs network was fixed. Were the mission of the organization, and so the needs network, allowed to alter an optimum of .825 would have been possible. Collectively, these results suggest that the organization was not particularly efficiently designed and/or there is substantial missing data about the organizational design. Further, the mission, as defined by the precedence ordering among tasks, limits what is possible in terms of congruence. The current organizational design requires 88 changes in who is doing what and has what resources in order to reach the optimal configuration. The hamming distance between the current and optimized design is indicative of the number of possible changes that can be made to move the organization closer to the optimum. If we treat this as a percentage, then 100 minus this number is the percentages of ties that will move the organization's design further from the optimum. The higher the hamming distance, the easier it will be, even by chance, to destabilize the organization. In this case, the hamming distance of 88 is 42% of the 208 possible linkages that could be changed. This indicates that a random change is slightly more likely to destabilize the organization and move it further from the optimum. We now take the original organization and ask, how should it be destabilized? Further, will these changes move it further from the optimum than a random change in the organizational design.

Locate vulnerabilities and select destabilization strategies.

Previous work has indicated that, from an adaptation perspective, a) node changes can be more devastating on system performance than relationship changes, and b) of the node changes those involving change in personnel are the most devastating (Carley, 2002c). We further argue that the removal or isolation of personnel is more practical, in the short term, than adding personnel, as the latter, particularly in covert networks, requires infiltration. For these reasons, we focus in this paper only on destabilization strategies associated with node removal.

In standard social network analysis, node changes are also the standard approach to network destabilization (Borgatti, 2002). Using standard social network techniques, individuals who are key in the social network are identified and then removed. The argument is that their removal serves to weaken or break the network so that messages flow slower and so that the network as a whole is no longer a single entity. There are several difficulties with this approach. First, since it only considers the social network it may be missing individuals who are critical due to what they are doing rather than who they know. Second, this approach assumes a static network – a single snapshot of who talks to or works with whom.

Table 3. Illustrative Measures of Criticality

Measure	Definition
Degree Centrality (Freeman, 1979)	Let M be the adjacency matrix representation of a square network. And let $n=\vert M\vert$. let $d_i = sum(M(i,:)) + sum(M(:,i))$ = out degree + in degree of node i Then Degree Centrality = d_i
Betweenness Centrality (Freeman, 1979)	Let G=(V,E) be the graph representation for the network. Fix a node $v \in V$. For any $(u,w) \in VxV$, let $n_G(u,w)$ be the number of shortest paths in G from u to w. If $(u,w) \in E$, then set $n_G(u,w)=1$. Define the following: let $S = \{(u,w) \in (V-\{v\})^2 \mid d_G(u,w) = d_G(u,v) + d_G(v,w)\}$ let $between = \sum_{(u,w)\in S} (n_G(u,v) * n_G(v,w)) / n_G(u,w)$ Then Betweenness Centrality of node v = between/(($\vert V\vert$-1)$*$($\vert V\vert$-2)), which normalizes the value to be in [0,1]
Cognitive Load (Carley. Lee and Krackhardt, 1999)	The Cognitive Load for agent i is defined as follows: let ATR = AT$*$RT’ let ATA = AT$*$AT’ let x_1 = # of agents that agent i interacts with / total # of agents = sum(A(i,:))/$\vert A\vert$ let x_2 = # of resources agent i manages / total # of resources = sum(AR(i,:))/$\vert R\vert$ let x_3 = # of tasks agent i is assigned to / total # of tasks = sum(AT(i,:))/$\vert T\vert$ let x_4 = sum of # resources agent i needs to do all its tasks / (total # tasks $*$ total # resources) = sum(ATR(i,:))/($\vert T\vert * \vert R\vert$) let x_5 = sum of # agents who do the same tasks as agent i / (total # tasks $*$ total # agents) = sum(ATA(i,:))/($\vert T\vert * \vert A\vert$) let x_6 = sum of negotiation needs agent i must do for each task / total possible negotiations = sum(AR(i,:) - ATR(i,:))/($\vert R\vert * \vert T\vert$) Then Cognitive Load for agent i = $(x_1+x_2+x_3+x_4+x_5+x_6)/6$
Task Exclusivity Ashworth and Carley, 2003	The Task Exclusivity Index (TEI) for agent i is defined as follows: $\sum_{j=1}^{\vert T\vert} AT(i,j) * e^{(1-sum(AT(:,j)))}$ The values are then normalized to be in [0,1] by dividing by the maximum TEI value.

In contrast, using dynamic network analysis we focus on criticality across the multiple matrices and across time. Nevertheless, for the sake of comparison, we contrast the removal of individuals identified as critical from both the standard and dynamic approach to networks. Typical measures of criticality for standard social networks are degree centrality and betweenness. To these we add cognitive load and knowledge/resource exclusivity. These measures are defined in Table 3.

These and other measures have been operationalized as C and C++ code. They are available as part of ORA and as part of NetStat at CMU. To determine which individuals ranked highest on each measure we ran ORA on the data set previously described. The results are shown in table 4. Note that two individuals are identified as critical. Moreover, the individual identified as critical via DNA measures (5, Ahmed) is actually quite low in the standard measures and would not have been picked up using a traditional approach.

Table 4. Individuals Identified as Critical in EB data set

Measure	Individual with Maximum Value	Type of Measure
Degree Centrality	7 Wadih al Hage	Standard Social Network
Betweenness Centrality	7 Wadih al Hage	Standard Social Network
Cognitive Load	5 Ahmed the German	Dynamic Social Network
Task Exclusivity	5 Ahmed the German	Dynamic Social Network

Determine performance characteristics in the short and long term after a destabilization strategy has been applied.

Given these measures, four distinct strategies for destabilizing the organization have been identified: eliminate the person with the highest degree centrality, betweenness centrality, cognitive load, or task exclusivity. We will measure the impact isolating the individuals high in these measures in two ways. First, we will contrast the relative resource congruence of the organizations without the isolated individual. This will be done in ORA. Second, we will contrast the relative change in performance in terms of accuracy and diffusion and ability to adapt to this change for the organization with and without these individuals. This will be done using DyNet. As a caveat, we note that although there are four criteria for identifying critical individuals, only two individuals are identified. Hence, for this data set, it is not possible to discriminate between the two different centrality measures. Nor is it possible to discriminate between the two dynamic social network measures.

The results of these removals are shown in table 5. All differences shown are significant. Neither removal substantially moves the design further from the optimal. Hence, we would expect the effects to be small. In addition, the removal of agent 5 actually increases resource congruence over the original design. On first blush, this is not good. However, keep in mind that resource congruence is a strict measure such that congruence is decreased when either agents do not have the resources needed for the task to which they are assigned or when agents have resources that are not necessary for the task that they are assigned. Removal of agent 5 is reducing the presence of unnecessary resources. Thus making the organizational design leaner. Making the organization optimal by reducing redundancy also make the organization less adaptive. Thus the removal of agent 5 makes the organization both more efficient but less adaptive.

In terms of performance the removal of agent 5 drops performance more than the removal of agent 7, and though it slows down the rate of recovery for performance, it has some what less of an impact than the removal of agent 7. Basically, think of performance as an S shaped curve. The removal of an agent can move the organization up or down on this curve depending on whether the individual was in a structurally disruptive position. But, since more/less change in performance is possible with the same degree of learning, some moves retard growth in performance more. In this case, the removal of agent 5 both lowers performance more and retards the growth more. From a performance standpoint, the removal of agent 5 should be more destabilizing and show a more prolonged effect.

If we explore diffusion the opposite is the case. For diffusion, the removal of agent 7 both lowers the initial diffusion more (compared to the removal of no agent) and it slows the rate at which diffusion is possible. Whereas, although the removal of agent 5 does drop the level of diffusion, it actually increases the rate of spread. In this case, the removal of agent 7 is more disruptive to the communication flows. It is important to keep in mind that this is the speed of information flow not the quality. Since the removal of agent 5 actually speeds the rate of information flow, it is speeding both the flow of accurate and inaccurate information. This potentially makes the organization more vulnerable to information warfare attacks.

Table 5. Impact of Agent Removal

Measure	Original Design	After Removal of 5	After Removal of 7
Hamming from Optimal	88	83	86
Resource congruence	.475	.525	.475
Performance as Accuracy – Initial Impact	78.5625	78.22	82.72
Performance Recovery – Percentage Increase in Performance	95.55	89.72	93.7
Diffusion - Initial	21.62291	14.70212	13.27369
Diffusion Recovery – Percentage Increase in Diffusion	71.23304	89.05325	50.87843

Limitations and Future Work

A key limitation of this work is that it is proceeding as though one had full information. In point of fact, that is unlikely to be the case. Hence, what is needed is a procedure for doing this analysis under varying levels of information assurance, placing confidence intervals around the results, and so placing the entire set of results in a decision context. A second limitation is that we have only illustrated the procedure with a very small data set. At issue is whether the measures identified are valuable for extremely large data sets. Work is proceeding on these and related issues. This work will help to define which measures are robust even under missing information and which measures are able to discriminate among nodes (potential critical individuals) in a meaningful way.

Herein we applied this procedure to an illustrative case. The results indicate that it would have been possible to disrupt the organization with the removal of a single agent. Moreover, the if the agent who is high in cognitive load or task exclusivity (critical in the overall meta-matrix) is removed the organization will be less adaptive, more efficient, exhibit lower performance, recover from the destabilization slower, but move information faster than the removal of an individual who was critical only in the social network. Sensitivity analyses, not shown, suggest that this is a robust result. Removal of high cognitive load individuals tends to be more disruptive than the removal of individuals high in degree centrality. Removal of either tends to disrupt the organization slightly more than a random removal. This latter effect may be increases if the organization is initially more optimally designed.

Future work should expand on this by considering other criteria for optimization, examining larger organizations where there are more complex networks, and explore other performance outcomes. Moreover, a key concern that needs to be addressed is the flow of incorrect information and the relative impact of such information warfare as opposed to personnel attacks.

References

[Borgatti, 2002] Borgatti, S.P., 2002, “The Key Player Problem,” Proceedings from National Academy of Sciences Workshop on Terrorism, Washington DC.

[Carley, 2002c] Carley, Kathleen M. 2002c, “Inhibiting Adaptation” In Proceedings of the 2002 Command and Control Research and Technology Symposium. Conference held in Naval Postgraduate School, Monterey, CA. Evidence Based Research, Vienna, VA.

[Carley, 2002b] Carley, Kathleen M. 2002b, “Dynamic Network Analysis” Paper presented at National Academy of Sciences/ National Research Council, Committee on Human Factors, Workshop on Dynamic Social Network Analysis, Washington D.C., November 2002. To be published in proceedings.

[Carley, 2002a] Carley, Kathleen M. 2002a, “Smart Agents and Organizations of the Future” The Handbook of New Media. Edited by Leah Lievrouw & Sonia Livingstone, Ch. 12 pp. 206-220, Thousand Oaks, CA, Sage.

[Carley, 1999] Carley, Kathleen M. 1999. “On the Evolution of Social and Organizational Networks.” In Vol. 16 special issue of Research in the Sociology of Organizations on “*Networks In and Around Organizations* “ edited by Steven B. Andrews and David Knoke. JAI Press, Inc. Stamford, CT, pp. 3-30.

[Carley, 1999] Carley, Kathleen M. & David Krackhardt, 1999, “A Typology for C^2 Measures.” *In Proceedings of the 1999 International Symposium on Command and Control Research and Technology*. Conference held in June, Newport,RI., Evidence Based Research, Vienna, VA.

[Carley, 1991] Carley, Kathleen M. 1991. “A Theory of Group Stability.” American Sociological Review 56(3): 331-354.

[Carley, Lee & Krackhardt, 1999] Carley, Kathleen M. Ju-Sung Lee and David Krackhardt, 2001, Destabilizing Networks, Connections 24(3):31-34.

[Carley & Svoboda, 1996] Carley, Kathleen M. & David M. Svoboda, 1996, "Modeling Organizational Adaptation as a Simulated Annealing Process." *Sociological Methods and Research*, 25(1): 138-168.

[Freeman, 1979] Freeman, L.C., 1979. Centrality in Social Networks I: Conceptual Clarification. Social Networks, 1, 215-239.

[Krackhardt & Carley, 1998] Krackhardt, D. and K. Carley, 1998, “A PCANS Model of Structure in Organization,” In: *Proceedings of the 1998 International Symposium on Command and Control Research and Technology*, 113-119.

[Lin & Carley, 2003] Lin, Zhiang and Kathleen M. Carley, 2003, Designing Stress Resistant Organizations: Computational Theorizing and Crisis Applications, Boston, MA: Kluwer.

[Ronfeldt & Arquilla, 2001] Ronfeldt, D. and J. Arquilla. September 21, 2001. "Networks, Netwars, and the Fight for the Future," First Monday, Issue 6 No. 10. online: http://firstmonday.org/issues/issue6_10/ronfeldt/index.htm.