

Kathleen M Carley, Ju-Sung Lee, David Krackhardt
2002 "Destabilizing Networks" in Connections 24(3): 79-92.

Destabilizing Networks¹

Kathleen M. Carley²

Ju-Sung Lee

David Krackhardt

Carnegie Mellon University, Pittsburgh, Pennsylvania, USA

The world we live in is a complex socio-technical system. Although social, organizational and policy analysts have long recognized that groups, organizations, institutions and the societies in which they are embedded are complex systems; it is only recently that we have had the tools for systematically thinking about, representing, modelling and analyzing these systems. These tools include multi-agent computer models and the body of statistical tools and measures in social networks.

This paper uses social network analysis and multi-agent models to discuss how to destabilize networks. In addition, we illustrate the potential difficulty in destabilizing networks that are large, distributed, and composed of individuals linked on a number of socio-demographic dimensions. The specific results herein are generated, and our ability to think through such systems is enhanced, by using a multi-agent network approach to complex systems. Such an illustration is particularly salient in light of the tragic events of September 11, 2001.

WHAT CAN OUR TOOLS DO?

There are a number of ways in which our tools, both classical social network techniques and the combination of networks and multi-agent systems, can help us understand network destabilization. Before describing these, an important word of caution is needed. Network tools are clearly not a panacea and it is important that as a community we do not oversell these tools. That being said, there are at least two fundamental ways in which network statistics and measures can be brought to bear to address issues at the heart of destabilizing networks.

¹This work was supported in part by the Office of Naval Research (ONR), United States Navy Grant No. N00014-97-1-0037, NSF IRI9633 662, Army Research Labs, NSF ITR/IM IIS-0081219, NSF KDI IIS-9980109, NSF IGERT: CASOS, and the Pennsylvania Infrastructure Technology Alliance, a partnership of Carnegie Mellon, Lehigh University, and the Commonwealth of Pennsylvania's Department of Economic and Community Development. Additional support was provided by ICES (the Institute for Complex Engineered Systems) and CASOS - the center for Computational Analysis of Social and Organizational Systems at Carnegie Mellon University (<http://www.casos.ece.cmu.edu>). The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Office of Naval Research, the National Science Foundation or the U.S. government.

²Direct all correspondence to: Prof. Kathleen M. Carley, Dept. of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA 15143, Email: kathelen.carley@cmu.edu, Fax: 1-412-268-6938, Tel: 1-412-268-3225, URL: <http://hss.cmu.edu/departments/sds/faculty/carley.html>

Location of critical individuals, groups, technologies

Given any network, such as a communication network, or alliance structure, or monetary flow, where the nodes are individuals, groups, computers, etc., a number of network measures such as centrality or cut-points can be used to locate critical nodes. Additional measures based on an information processing view of organizations also exist for locating critical employees, redundancy, and potential weak points within groups and organizations. Many of the traditional social network measures and the information processing network measures are embedded within ThreatFinder (Carley, 2000). ThreatFinder is a computer program that uses a combination of network analysis and multi-agent modelling to determine the potential information security risk from personnel that an organization faces due to its architecture. The degree, type, and location of possible threats, such as critical employees and lack of redundancy are assessed. These "location" techniques are useful within companies to help ensure information security and are useful within and among groups and organizations in mitigating the effectiveness of networks. For example, individuals or groups with the following characteristics can be identified:

1. An individual or group where removal would alter the network significantly; e.g., by making it less able to adapt, by reducing performance, or by inhibiting the flow of information. Illustrative nodes are those high exceptionally high in centrality (Bonacich, 1987) or high in structural holes (Burt, 1992).
2. An individual or group that is unlikely to act even if given alternative information. This can be found as an individual high in centrality and Simmelian ties (Krackhardt, 1999).
3. An individual or group that if given new information can propagate it rapidly. Such individuals may be seen as gossips, innovators, or early adopters (Rogers and Shoemaker, 1971). Possible indicators are high degree centrality or high structural holes.
4. An individual or group that has relatively more power and can be a possible source of trouble, potential dissidents, or potential innovators. Individuals with relatively more power may be high in centrality (Bonacich, 1987; Brass, 1991; Brass and Burkhardt, 1992). Possible innovators may be those who are isolates or those who have moved about so much that they have broad and distributed knowledge and contacts.
5. An individual or group where movement to a competing group or organization would ensure that the competing unit would learn all the core or critical information in the original group or organization (inevitable disclosure) (Carley, 2000).
6. An individual, group, or resource that provides redundancy in the network (Carley and Ren, 2001). Measures of redundancy are available in ThreatFinder (Carley, 2000).

For the measures discussed above most can be calculated using UCINET³ or the meta-network R-package package⁴.

Pattern location

Over the past few years, major advances have been made in graph level analysis. These techniques include the P* family of tools, network level metrics (such as group and graph clustering algorithms using distance metrics such as the Hamming distance). These pattern location techniques can be used on any data that can be represented as graphs; such as, interaction or communication networks, monetary networks, inter-organizational alliances, mental models, texts, web pages, who was present at what event, and story lines. These pattern location techniques, particularly when combined with machine learning techniques, are likely to be especially powerful for locating patterns not visible to the human eye. A key to many of the detection algorithms is that they search for behavior that is different

³ <http://eclectic.ss.uci.edu/~lin/ucinet.html>

⁴ <http://legba.hss.cmu.edu/R.stuff>

from some baseline. Thus, if run on network data, The baseline might be networks, biased networks, or a sample of existing networks. For example, the following kinds of patterns or breaks in patterns can be examined:

- The basic components that account for the networks structure can be identified; e.g., the number and types of sub-groups, or the number of triads, stars, and the extent of reciprocity (Anderson, Wasserman, and Crouch, 1999; Wasserman, and Pattison, 1996).
- The central tendency within a set of networks, and the networks that are anomalous when contrasted with the other networks can be located (Banks and Carley, 1994).
- Critical differences between two or more sets of networks can be identified; e.g., are programming teams structured differently than sales teams or are managers' mental models different from subordinates (Banks and Carley, 1994; Carley and Banks, 1993; Butts and Carley, 2001). For sets of concepts, comparison techniques based on the idea of lossy integration and set theory have been used to compare two or more concept networks or mental models (Carley and Palmquist, 1992; Carley, 1997). In principle, these methods developed for text analysis could be utilized for the comparison of social networks.
- Which components in the network are structured significantly differently from the rest of the overall network? A standard approach is to locate the nodes or sets of nodes that differ significantly from other nodes on standard measures such as degree centrality, betweenness, and number of cliques. However, for extremely large networks or where only samples of data on the network exist this approach may not be feasible (processing time is excessive, space requirements are too high, or missing data is too high). Under these conditions, you can use machine learning algorithms such as simulated annealing (Kirkpatrick, Gelato and Vichy, 1983) or Bayesian updating (Butts, forthcoming; German, Carlin, Stern, and Rubin, 1995; Robert, 1994) to search through the network to locate the node or set of nodes that are highest on some criteria or best match some criteria such as excessively high or low centrality.
- Whether the existing network is coherent; i.e., what is the likelihood that there are key missing nodes or relations. One approach here is to locate the differences between an actual network and a network predicted from first principles to see where there are differences. For example, if two individuals are not interacting in the social network but should be based on the principles of relative similarity and relative expertise, then there may be hidden relations. This is one of the calculations in ThreatFinder (Carley, 2000).

What-if analysis and policy guidance

In addition, multi-agent models of adaptive agents embedded in social networks can be used to address issues of network destabilization by providing managerial and policy guidance (Carley, forthcoming a). In a multi-agent computational program the behavior of the group or organization emerges from the actions and interactions of the agents who are members of the group or organization. Typically the agents are able to learn and adapt, although models vary widely in the extent to which the agents are cognitively realistic (Carley, forthcoming b). Few multi-agent models have more than 100,000 agents and in general the number of agents decreases as the cognitive complexity and realism of the agents increases. Multi-agent systems are typically non-linear and exhibit path dependence. Most multi-agent models have no network underpinning. In the artificial life models (Epstein and Axtell, 1997) the agents typically interact on a grid with physical proximity serving as a proxy for networks. In the most cognitively sophisticated models, such as the Soar models (Tambe, 1997), the set of interactions and so the network are predefined. However, recently, there has been a movement to combining multi-agent and network models (More and Ramanujam, 1999; Levinthal, 1997; Macy and Skvoretz, 1998; Carley, 1990; Carley and Svoboda, 1997).

Multi-agent network models, if based on known information about general or specific characteristics of groups, can suggest general or specific guidance about how to affect or protect the underlying group, organization or society. Exactly what these models can address depends on the purpose of the model and its veridicality. Following is a series of illustrative examples of potential applications where various researchers using multi-agent network models have worked or are working:

- Suggesting factors that make groups adaptive or maladaptive (Carley and Lee, 1998).
- Examining the efficacy of different policies for destabilizing networks; e.g., what kinds of networks can be destabilized by simply removing the leader (Arquilla and Ronfeldt, 2001)? What are the characteristics of networks that are difficult to destabilize (Watts, 1999; Carley, forthcoming a)?
- Examining the efficacy of different data collection and privacy policies. For example, would we be more likely to mitigate a bioterrorist attack if we kept absentee data or if we tracked hits on web based medical information pages (Carley, Yahja and Fridsma, 2001)?
- Predicting the rate of information diffusion and the impact of different technologies for spreading information and so changing beliefs through social influence processes (Oram, 2001; Watts, 1999; Carley, forthcoming c; Macy and Strang, forthcoming).
- Predicting voting outcomes or likelihood of consensus in groups, given the existing social networks and initial beliefs (Friedkin, 1998; Bueno De Mesquita and Stokman, 1994).
- Suggesting factors that can slow the rate of response by a network to a new situation or event, mitigate the emergence of new behaviors, and limit the ability of the network to adapt (Wegner 1995; Axtell, 2000; Carley, forthcoming a).
- Predicting civil violence (Epstein, Steinbrunner and Parker, 2001)
- Determining how close your group or company is to having its core competencies and processes discovered by another group (i.e. inevitable disclosure) (Carley, 2000).
- Examine the efficacy of different marketing and information warfare strategies (Pew and Mavavor, 1998, ch. 11).

Doubtless each researcher in this area has thought of these and other possible applications. We note that at the moment there are a number of difficulties in applying existing tools to complex socio-technical systems. First, most of the existing multi-agent network models are implemented for small networks. Even when the underlying measure can be used on large networks, containing 1000s or 10,000s of nodes, the underlying computer software or hardware often limits the feasible analysis to small networks, those less than a few hundred nodes. For example, UCINET can handle large node sets, but, in practice the memory limitations on the machine on which it is run and the lack of parallelization procedures means that it is an impractical tool for networks of tens of thousands of nodes. Second, we have no public databases of large networks on which to test new technologies. However, large networks based on web linkages are being developed. Third, the existing measures and tools work best when the data is complete, i.e., when we have full information about the links among the nodes. However, large scale distributed networks may have considerable missing data. We will at best have sampled information, some of the information may be intentionally hidden (hence missing data may not be randomly distributed), the data is likely to be at different time scales and layers of granularity, and the cost and time to get complete information may be prohibitive. Thus, we need to begin to address issues of sampling, of estimating the impact of missing information, of estimating networks given basic human cognitive properties and population level and cultural data, and in combining data from alternative and dispersed sources using techniques such as multiple imputation (Rubin, 1987, 1996; Schafer, 1997; Yuan, 1990). There are obviously other difficulties, but even these provide some guidance for what to expect when applying our existing tools to complex socio-technical systems.

WHY MIGHT IT BE DIFFICULT TO DESTABILIZE DISTRIBUTED NETWORKS?

One possible approach at overcoming, or at least ameliorating, some of these difficulties is to use computational analysis, where the models combine multiple cognitively realistic agents and social networks. We now illustrate the use of such models to address the issue of network destabilization. As noted, socio-technical systems are complex. First, let us consider the source of complexity. We can point to a large number of sources of complexity: e.g., new technologies, emergent cultures, complex trade laws, etc. At a more fundamental level there are two very dominant sources: (1) humans adapt and (2) humans interact. Humans adapt in part because they can learn, but what they learn is limited because they are boundedly rational. Human interactions are of course influenced by the web of affiliations (kinship, religion, economics, etc.) that interlock people to varying degrees at different times. Since individuals can adapt and are woven together into a complex network, the groups, organizations and institutions of which they are members also have these properties. Thus, we have intelligent adaptive agents and multiple networks. However, these are not de-coupled systems. Humans learn when they interact with each other and what they learn changes the knowledge network (who knows what), with whom they interact (the social network), and how they perform tasks. Who you know and what you know are linked together in a feedback loop. The result is that the networks in which people are embedded are dynamic.

Network dynamics is a function of not just the social network, but a meta-matrix of networks – not the least of which are the knowledge network (who knows what), the information network (what ideas are related to what), and the assignment network (who is doing what) (Carley and Hill, 2001, Krackhardt and Carley, 1998). A highly simplified version of this meta-matrix representation of the meta-network is shown in Table 1, where for the sake of simplicity only the networks related to agents, knowledge and tasks are shown. As noted by Agranoff and McGuire (1999) “the ability to tap the skills, knowledge, and resources of others is a critical component of networking capacity,” the ability to manage the organization. Similarly, to determine how to change or destabilize a network, then, it is important to consider the further webs in which a social network is situated and the way in which human cognition operates (Krackhardt, 1990; Carley and Hill, 2001).

| | Agents | Knowledge | Tasks |
|-----------|----------------|---------------------|-------------------------|
| Agents | Social Network | Knowledge Network | Assignment Network |
| Knowledge | | Information Network | Needs Network |
| Tasks | | | Task-Precedence Network |

We have built a relatively simple computational model of this dynamic process — CONSTRUCT-O (for a description of this model, see Carley and Hill, 2001). Such models are valuable in addressing theoretical, social, managerial and policy issues (Carley, 2001; Carley and Gasser, 1999; Epstein and Axtell, 1997). A key feature of these models is that they let us think systematically about the ramifications of policies, at a scale not comprehensible by the unassisted human mind, and so can help uncover major problems. We can use this model to address the question “what leads to the destabilization of networks?” It is worth noting that the predecessor of this model, CONSTRUCT, was used to examine the factors enabling group stability (Carley, 1990; 1991) and the evolution of networks (Carley, 1999).

The model works by first assuming a set of agents who differ in terms of their socio-demographic characteristics (such as age, gender, education), their knowledge and beliefs. Individuals also forget. Individuals interact if they are available for interaction and are motivated to do so. There are two basic motivations to interact – relative similarity and relative expertise – both of which are basic to human nature. Relative similarity is the tendency of people to choose to interact with those who are more

similar. Relative expertise is the tendency of people to seek out new information from those whom they perceive to be more expert. When people interact they learn and their learning changes whom they view as relatively similar or expert, how well they perform the tasks to which they are assigned, and who can be assigned to which tasks.

These changes also alter whether or not there is an emergent leader and which individual takes on that role (Cohen, Bennis and Wolkon, 1962). Individuals are more likely to develop effective leadership skills if they have high cognitive ability, prior experience (Atwater, Dionne and Avolio, 1999), and extroversion (Kickul and Neuman, 2000). Individuals who have high cognitive ability and experience typically take on more tasks, are given more resources, and have more knowledge. Prior experience and extroversion often lead to a wider range of interaction partners. Stress typically occurs when cognitive load increases. Additionally, individuals are likely to emerge as leaders if they have high stress tolerance, have strong self-esteem (Atwater, Dionne and Avolio, 1999) and are open to new experiences (Kickul and Neuman, 2000). As such they are likely to be willing to tell others what to do, shed tasks, give away resources, etc. Individuals with high cognitive loads are likely to be emergent leaders for a variety of reasons including they are most likely to tell others to do things (i.e., shed tasks) and most likely to be in a position of power in terms of what and whom they know. An agent is more likely to be an emergent leader and to direct the activity of the distributed network, even if only temporarily, if that agent is in a strong structural position in the social, knowledge and assignment networks. Overall cognitive load, not simply structural power, is key to tracking who is likely to be the emergent leader. Based on these considerations, we define the emergent leader as the individual with the highest cognitive load (the most people to talk to, the most information to process, the most tasks to do, the hardest tasks to do, the most people to negotiate with to get the job done, etc.) (Carley and Ren, 2001).

The cognitive resources of the group and the leader, the cognitive load, and the behavior of the leader have a combined impact on performance (Fiedler, 1986). Consequently, emergent leaders, by virtue of their centrality across the entire meta-network are good candidate agents to remove if the goal is to destabilize the network. Therefore, the effect of node extraction on network evolution will be examined by removing the emergent leaders from the networks at a particular point in time and then seeing how the networks evolve.




Figure 1. A Stylized Hierarchical Centralized Network

There are at least three indicators of destabilization. One is where the rate of information flow through the network has been seriously reduced, possibly to zero. A second is that the network, as a decision-making body, can no longer reach consensus, or takes much longer to do so. A third is that the network, as an organization, is less effective; e.g., its accuracy at doing tasks or interpreting information has been impaired. There are other instances of network instability, but such measures are sufficient for this brief introduction.

linking individuals to knowledge/resources, tasks are needed, as are the networks linking knowledge to tasks and tasks to tasks. For example, in Figure 1, the most central agent, although interacting with the most others and being cognitively more similar to the most others (closeness in physical space), is not the emergent leader. The reason is that this social network is linked into a set of networks denoting who knows what, does what, what is needed to what tasks, the order in which tasks need to be done, and so on.

When they visually examine the hierarchical network, most people will predict that removal of either the leader or the central agent will be most likely to destabilize the structure. Further, given just the social network, most people predict that the most central agent is likely to emerge as the new leader. In contrast, for the distributed decentralized structure, Figure 2, it is not clear whether there is a single node that could be removed to destabilize the network. There is substantial disagreement among people who examine this network over which node to remove to destabilize the network, and even over whether it is even possible to destabilize the network. This is the case even when, as in Figure 2, the emergent leader is the most central agent. Further, there is little agreement over who will emerge as the leader.

To really determine whether removal of a node will destabilize a structure we need to account for adaptation. Since individuals can learn, the underlying social networks are dynamic. They will change whether or not various nodes are removed. Further, individual learning will lead the overall structure to adapt, often in unforeseen ways as nodes are removed or isolated. As a result, removing a node may result in a new emergent leader. This new emergent leader cannot be predicted just from the social network. A possible path of change for the hierarchical network in Figure 1 is shown in Figure 3 and a transition path for the distributed network of Figure 2 is shown in Figure 4. In each graph, the emergent leader is again shown as a rectangular node labeled 'LEADER' and the most central agent as an oval node labeled 'CENTRAL.' In addition, to help orient the reader, when an agent is removed the position that that agent would have had if he/she had not been removed is labeled with the word 'REMOVED.'

REMOVED




Figure 3a. Removal of an Emergent Leader in a Stylized Hierarchical Centralized Network - Immediate Response to Removal of Emergent Leader

For the hierarchy, we begin with the hierarchy shown in Figure 1. Initially, the emergent leader's cognitive load is significantly higher than the subordinates in the hierarchy. Then over the course of the simulation the emergent leader is extracted. Upon destabilization, the distribution of cognitive load shifts such that more agents have higher loads, and more than one leader emerges. Figure 3a contains the resultant network that emerges after the original emergent leader is removed. Immediately, the extraction of the leader agent in Figure 1 causes the hierarchy to break up into two smaller networks.

Using this model we examine two very distinct structures – a hierarchical centralized structure and a distributed decentralized one. For both structures, although different in scale, the underlying distributions of knowledge/resources and tasks are similar as are the networks linking knowledge/resources to tasks and tasks to tasks. These other networks are not shown as the figure becomes unwieldy; however, they do impact who learns what over time and so changes in the social network and cognitive load. The Krackplot representations of only the social network component of these structures are displayed in Figures 1 (hierarchical) and 2 (decentralized). In Figures 1 and 2, the spatial arrangement of nodes represents knowledge proximities between agents (i.e., the closer two nodes the more likely they have similar knowledge). Those closer together also tend to share more knowledge. The amount of knowledge, resources and tasks associated with each individual agent is not shown. Individuals seek out others who (1) are similar, knowledge-wise and (2) can provide the resources for completing his or her tasks. A line connecting two agents indicates that during the window of observation these two agents interacted with each other. The bold-lines denote strong interaction network ties that occur when an agent has established a relationship that is part functional (i.e., task-resource based) and part social (i.e., general knowledge and demographic based).

A rectangular node labeled 'LEADER' denotes the "Emergent Leader" agent. This agent is the individual with the highest cognitive load (i.e., most resources, tasks, and communication/network ties). An oval node labeled 'CENTRAL' denotes the agent with the most network ties. If the agent is both the emergent leader and the most central then a rectangular node labeled 'LEADER/CENTRAL' denotes that agent. Some agents may share information with others but are nevertheless not interacting with any of the other agents during a particular window of observation. Such agents will appear as isolated nodes with no lines connecting them to other agents.




Figure 2. A Stylized Distributed Decentralized Network

It is important to note that if you only observe the social network, as in Figures 1 and 2, you cannot determine who has the highest cognitive load and is therefore likely to emerge as a leader. In the hierarchal network (Figure 1), or for that matter in any network, the emergent leader is not necessarily the most central agent. If we were to only look at the social network, we might assign leadership on the basis of the power of the agent's structural position. That is, examining just the interaction matrix one might be tempted to conclude that the agent with the highest degree centrality or betweenness was the leader. However, this can be misleading. While there is often a correlation between an agent's position in the social network and their overall cognitive load, it is not perfect. Centrality is only one of the factors that enters into the overall calculation of a cognitive load. To determine loads, the networks

Once the leader is extracted the network reforms with two emergent leaders who are essentially competing for control – neither of which is the most central agent. After further simulation, the network has adapted to the loss and a new single leader has emerged (see Figure 3b). In reforming itself back into a hierarchy, a new leader emerges whose cognitive load is higher than that of the first leader, indicative of a less pure hierarchy. Not all hierarchies will change in this way – but this diagram is illustrative of the impact of extracting a leader on a hierarchical network.

Removing the leader in a hierarchy not only destabilized the network, it also makes the overall communication structure more decentralized. When centralized groups become decentralized initial leaders are often demoted and moved to positions of least importance (Cohen, Bennis, and Wolken, 1962). Cohen, Bennis and Wolken (1962) suggested that such a change may be a psychological response to imposed leadership. Our analysis suggests that this may simply be the result of structural differences in the meta-network which lead to differences in cognitive load. Notice that the leader on the right in Figure 3a is demoted in 3b.




Figure 3b. Removal of an Emergent Leader in a Stylized Hierarchical Centralized Network – Eventual Response

This example illustrates that destabilizing a hierarchy may have unintended consequences — demotion of leaders and initial in-fighting. It also illustrates that visual inspection of the social network alone led to an incorrect prediction as to who would emerge as the new leader. We might ask, what if the central agent rather than the leader was removed. Further simulation analysis shows that not only does the hierarchy not break into factions initially, but its performance is hardly even affected. For hierarchies, the simulation analysis suggests that regardless of the size of the hierarchy, removal of the leader degrades performance more than removal of the central agent. Moreover, hierarchies, relatively quickly restabilize with only a single new emergent leader. A number of actions may have consequences similar to node removal: e.g., isolating, hiring away the leader, reducing the number or complexity of tasks the leader is doing, or stopping the flow of information or resources through all links connected to the leader. For the hierarchical network, the leader's ability to control the hierarchy can also be decreased by adding new links in the social network. Such additional linkages can also lead to performance drops.

In Figure 4, the consequences of removing an emergent leader on a distributed decentralized network are portrayed. The initial structure is that in Figure 2. As with the hierarchy, during the course of the simulation the emergent leader, LEADER/CENTRAL, is now extracted. In Figure 4a, like Figure 3a, the position that the original leader would have held if he/she had not been extracted is denoted by the word 'REMOVED'. In Figure 4a we see that after that a new leader emerges in the same vicinity as the

original LEADER. However, this newly emergent leader is neither the most central nor does he/she re-establish the ties that were lost with the former leader. In the long run, Figure 4b, multiple new leaders emerge. In addition, the agent who in Figure 4a was the most central also becomes an emergent leader. A third leader emerges in a structural position very similar to that of the original leader (who was removed). The fact that two of the new leaders are near the original leader is indicative of the fact that the structure of the task, knowledge and resource networks (which are not visible) in that vicinity promotes the development of emergent leaders. Further, when the original leader was present, that agent was inhibiting the emergence of alternative leaders. The original leader had maintained key resources, knowledge and important ties. The original leader had played the role of the gatekeeper between the left and right sides of the network. Once the agent LEADER/CENTRAL was removed, tasks and resources could be redistributed, agents had to rely on other experts, and multiple leaders could eventually emerge.




Figure 4a. Removal of an Emergent Leader in a Stylized Distributed Decentralized Network - Immediate Response to Removal of Emergent Leader

Computational analysis reveals that even the removal of the LEADER/ CENTRAL agent may have unforeseen effects. In the distributed network adding or dropping links is as likely to increase an individual node's power as to decrease it. Consequently, the overall impact of removing the leader in a distributed network is not as likely to create a power vacuum as in the hierarchical network. If this is the case, then removal of that agent will have little impact. It may be necessary to simultaneously remove more nodes to have the same impact on a distributed decentralized system as removing one node would have on a hierarchy. In this sense, the problem of destabilization is more difficult for a distributed than for a hierarchical network. We might ask what if the leader was not also central. As with the hierarchy, further simulation reveals that the removal of the central agents as opposed to the leaders is less likely to degrade performance. Computational analysis also reveals that removal of a single node does not transform the structure, despite agent adaptation; i.e., hierarchies remain as hierarchies and distributed structures remain distributed.

We note that many resistance groups are organized as distributed decentralized networks. For example, in the Earth Liberation Front (ELF) according to ELF publicist, Craig Rosebraugh, there is a "series of cells across the country with no chain of command" (Barr and Baker, 2001). In such cases, there is "no central leadership where they can go and knock off the top guy and it will be defunct" (Barr and Baker, 2001). Our analysis suggests further that even if you find emergent leaders, removing them

simply paves the way for new leaders to emerge and the overall network will remain more or less intact. However, unlike the hierarchy, the removal of the initial LEADER may serve to, in the long run, increase internal fighting as multiple LEADERS are likely to eventually emerge. The hierarchy splits in to factions then reforms as a hierarchy with one leader, the distributed system does not faction at first but may eventually as multiple leaders emerge.

To really track and understand network dynamics, to really be able to determine how to destabilize networks, we need to consider the position of individuals and groups as they are embedded in the overall meta-network. We need to move beyond embeddedness in the social network (Granovetter, 1985) to overall embeddedness in the meta-network. Although he does not use the network nomenclature, this is essentially Schein's (1985) point in his discussion of leadership.




Figure 4b. Removal of an Emergent Leader in a Stylized Distributed Decentralized Network - Eventual Response

Herein we used cognitive load to track embeddedness in the overall meta-network linking personnel, knowledge/resources and tasks. Now examine the change in the distribution of cognitive load for the distributed decentralized network (Figure 5). These distributions, going from left to right, correspond to Figure 2, Figure 4a and 4b respectively. The original leader has a much higher cognitive load than do other members of the distributed decentralized structure. Initial destabilization results in multiple

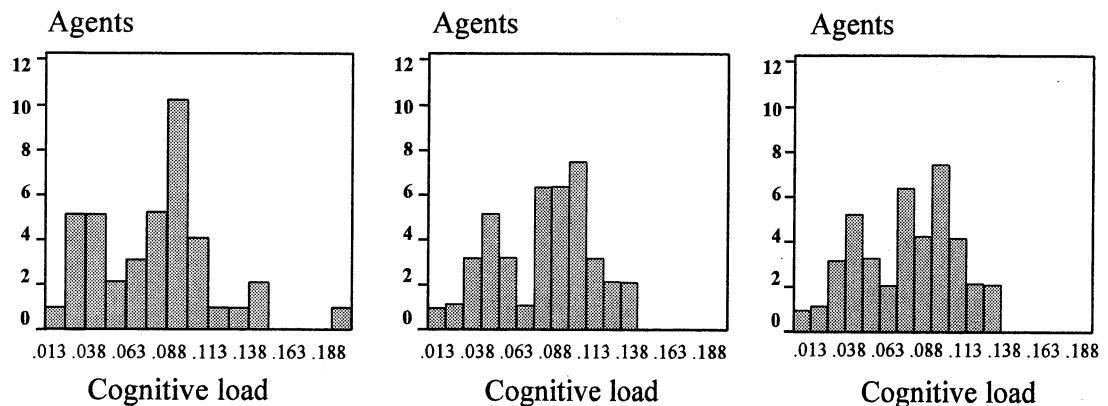


Figure 5. Change in the distribution of Cognitive load

emergent leaders forming, as indicated by the subsequent distributions of cognitive load. While the number of emergent leaders drops as the network re-stabilizes, the emergent leaders are not as distinctive as the original.

As network theorists, we often think about networks as snapshots – pictures of a group at a point in time. The techniques and tools that have been developed over the past several decades are extremely useful in understanding such networks (assuming of course that the data is complete or almost so). Moreover, we often think of networks primarily in terms of a relatively small, single relation and single type of node; e.g., friendship among students. At this point in time, few tools are available to the analyst interested in large, adaptive, multi-plexed, multi-coloured networks with high levels of missing data. The development of such tools is necessary if we are to successfully meet the challenge of understanding, predicting and explaining the behaviour of multi-agent networks of this ilk. Whether the topic is terrorism, the global economy or the nature of the Internet, we are dealing with complex socio-technical systems that are large, multiplex, multi-nodal and adaptive. It is critical that we rise to this challenge and develop a new set of tools combining the methodologies of social networks and computer science. Without such tools, we will be theorizing in the dark.

REFERENCES

- Agranoff, R. and M. McGuire. 1999. Managing in Network Settings, *Policy Studies Review*, 16 (1): 18-41.
- Anderson, C., S. Wasserman, and B. Crouch. 1999. A p* Primer: Logit Models for Social Networks *Social Networks*, 21: 37-66.
- Arquilla, J. and D. Ronfeldt (editors). 2001. Networks and Netwars: The Future of Terror, Crime, and Militancy. Santa Monica, Calif.: RAND, MR-1382-OSD. www.rand.org/publications/MR/MR1382/
- Atwater, L.E., S.D. Dionne, and B. Avolio. 1999. A longitudinal study of the leadership development process: individual differences predicting leader effectiveness. *Human Relations* 52 (2): 1543-62.
- Axtell, R. L. 2000. Effects of Interaction Topology and Activation Regime in Several Multi-Agent System, Brookings Working Paper, <http://www.brook.edu/ES/dynamics/papers/interaction/>.
- Banks, D. and K.M. Carley. 1994. Metric Inference for Social Networks. *Journal of Classification*, 11: 121-149.
- Barr, D. and A. Baker. January 8, 2001. Getting the Message from 'Eco-Terrorists': Mystery Group Takes Its Campaign East. *New York Times*, p. A15.
- Bonacich, P. 1987. Power and Centrality: A Family of Measures. *American Journal of Sociology*, 92: 1170-82.
- Brass D.J. 1991. Power in Organizations: A Social Network Perspective. In: *Research in Politics & Society*, G. Moore and J.A. Whitt, Eds. Stamford, CT: JAI Press.
- Brass D.J. and M.E. Burkhardt. 1992. Centrality and Power in Organizations. In: *Networks and Organizations*, Nohria, Nitin and R. Eccles, Eds. Cambridge, MA: Harvard Business School Press.
- Bueno De M., B. Stokman, and F. Stokman. 1994 *European Community Decision Making: Models, Applications, and Comparisons*, New Haven: Yale University Press.
- Burt, R. 1992. *Structural Holes: The Social Structure of Competition* Harvard University Press: Cambridge, MA.
- Butts, C. T. and K.M. Carley. 2001. Multivariate Methods for Inter-Structural Analysis. CASOS Working Paper, http://www.casos.ece.cmu.edu/working_papers.html, Carnegie Mellon University.
- Butts, C.T. Forthcoming. Network Inference, Error, and Informant (In)accuracy: A Bayesian Approach, *Social Networks*.

- Carley, K.M. Forthcoming a. Computational Organization Science: A New Frontier . *Arthur M. Sackler Colloquium Series on Adaptive Agents, Intelligence and Emergent Human Organization: Capturing Complexity through Agent-Based Modeling* October 4-6, 2001; Irvine, CA, National Academy of Sciences Press.
- Carley, K.M. Forthcoming b. Intra-Organizational Computation and Complexity. In: *Companion to Organizations*, J.A.C. Baum, Ed., Blackwell Publishers.
- Carley, K.M. Forthcoming c. Smart Agents and Organizations of the Future. In: *Handbook of New Media*, L. Lieveroux and S. Livingstone, Eds.
- Carley, K.M. 2001. Computational Organization Science: A New Frontier. Paper prepared for the NAS sponsored Sackler Colloquium, Irvine CA, October 2001.
- Carley, K.M. 2000. Information Security: The Human Perspective. CASOS Working Paper, Carnegie Mellon University, http://www.casos.ece.cmu.edu/working_papers.html.
- Carley, K.M. 1999. On the Evolution of Social and Organizational Networks. In Steven B. Andrews and David Knoke (Eds.) Vol. 16 special issue of *Research in the Sociology of Organizations*. on "Networks In and Around Organizations." JAI Press, Inc. Stamford, CT, pp. 3-30.
- Carley, K.M. 1997. Extracting Team Mental Models Through Textual Analysis. *Journal of Organizational Behavior*, 18: 533-538.
- Carley, K.M. 1991. A Theory of Group Stability. *American Sociological Review*, 56(3): 331-354.
- Carley, K.M. 1990. Group Stability: A Socio-Cognitive Approach. In: *Advances in Group Processes: Theory & Research. Vol. VII*. E. Lawler, B. Markovsky, C. Ridgeway and H. Walker, Eds., 1-44. Greenwich, CN: JAI Press.
- Carley, K.M. and D. Banks. 1993. Nonparametric Inference for Network Data. *Journal of Mathematical Sociology*, 18(1): 1-26.
- Carley, K.M. and L.S. Gasser. 1999. Computational Organization Theory. In: *Distributed Artificial Intelligence*. Cambridge, G. Weiss (Ed MA: MIT Press.
- Carley, K.M. and V. Hill. 2001. Structural Change and Learning Within Organizations. In: *Dynamics of Organizations: Computational Modeling and Organizational Theories*, A. Lomi and E.R. Larsen, Eds., 63-92. MIT Press/AAAI Press/Live Oak.
- Carley, K.M. and J-S. Lee. 1998. Dynamic Organizations: Organizational Adaptation in a Changing Environment. In: *Advances in Strategic Management, Vol. 15, Disciplinary Roots of Strategic Management Research*, J. Baum, Ed., 269-297. JAI Press.
- Carley, K.M. and M. Palmquist. 1992. Extracting, Representing and Analyzing Mental Models . *Social Forces*, 70(3): 601-636.
- Carley, K.M. and Y. Ren. 2001. Tradeoffs Between Performance and Adaptability for C³I Architectures. In: *Proceedings of the 2001 Command and Control Research and Technology Symposium*, Annapolis, Maryland, June, 2001.
- Carley, K.M. and D.M. Svoboda. 1996. Modeling Organizational Adaptation as a Simulated Annealing Process." *Sociological Methods and Research*, 25(1): 138-168.
- Carley, K.M., A. Yahja, and D. Fridsma. 2001. BioWar. CASOS online report, CMU, Pittsburgh, PA <http://www.casos.ece.cmu.edu/biowar/index.html>
- Cohen, A.M., W.G. Bennis, G.H. Wolken. 1962. The Effects of Changes in Communication Networks on the Behaviors of Problem-Solving Groups. *Sociometry*, 25(2): 177-196.
- Epstein, J.M. and R. Axtell. 1997. *Growing Artificial Societies*, Boston, MA: MIT Press.
- Epstein, J.M., J.D. Steinbrunner, and M.T. Parker. 2001. Civil Violence: An Agent-Based Computational Approach, Brookings Working Paper, <http://www.brook.edu/es/dynamics/papers/cviolence/cviolence.htm>

- Fiedler, F.E. 1986. The contribution of cognitive resources and leader behavior to organizational performance. *Journal of Applied Social Psychology*, 16(6): 532-48.
- Friedkin, N. 1998. *A Structural Theory of Social Influence*, Cambridge University Press.
- Gelman, A., J.B. Carlin, H.S. Stern, and D.B. Rubin. 1995. *Bayesian Data Analysis*. London: Chapman and Hall.
- Granovetter, M. 1985. Economic Action and Social Structure: The Problem of Embeddedness. *American Journal of Sociology*, 91: 481-510
- Kickul, J. and G. Neuman. 2000. Emergent Leadership Behaviors: The Function of Personality and Cognitive Ability in Determining Teamwork Performance and KSAs. *Journal of Business and Psychology*, 15(1): 27-51.
- Kirkpatrick, S., C.D. Gelatt and M.P. Vecchi. 1983. Optimization by Simulated Annealing. *Science* 220 (4598): 671-680.
- Krackhardt D. 1990. Assessing the Political Landscape: Structure, Cognition, and Power in Organizations. *Administrative Science Quarterly*, 92, 142-154.
- Krackhardt, D. 1999. The Ties that Torture: Simmelian Tie Analysis of Organizations. *Research in the Sociology of Organizations*, 16: 183-210.
- Krackhardt, D. and K.M. 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. June. Monterey, CA. Levinthal, D. 1997. Adaptation on rugged landscapes. *Management Science*, 43: 934-950.
- Macy, M. and J. Skvoretz. 1998. The Evolution of Trust and Cooperation Between Strangers: A Computational Model. *American Sociological Review*, 63: 638-660.
- Macy, M. and D. Strang. forthcoming. 'In Search of Excellence: Fads, Success Stories, and Adaptive Emulation. *American Journal of Sociology*.
- Oram, A., Ed. 2001. *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*. Sebastopol, Calif.: O'Reilly and Associates.
- Pew, R.W. and A.S. Mavavor, Eds. 1998. *Modeling Human and Organizational Behavior: Application to Military Simulations*, Final Report, National Research Council, National Academy Press. Washington D.C.
- Rogers, E.M. and F.F. Shoemaker. 1971. *Communication of Innovations: A Cross-Cultural Approach*, (Second edition). New York: Free Press.
- Robert, C.P. 1994. *The Bayesian Choice: A Decision-Theoretic Motivation*. New York: Springer.
- Rubin, D.B. 1987. *Multiple Imputation for Nonresponse in Surveys*. New York: John Wiley and Sons, Inc.
- Rubin, D.B. 1996. Multiple Imputation After 18+ Years. *Journal of the American Statistical Association*, 91: 473-489.
- Schafer, J.L. 1997. *Analysis of Incomplete Multivariate Data*. Boca Raton, FL: Chapman & Hall/CRC.
- Schein, E. 1985. *Organizational Culture and Leadership*. San Francisco: Jossey-Bass.
- Wasserman, S. and P. Pattison. 1996. Logit Models and Logistic Regressions for Social Networks, I: An introduction to Markov Graphs and p^* . *Psychometrika*, 60: 401-426.
- Watts, D. J. 1999. *Small Worlds: The Dynamics of Networks Between Order and Randomness*. Princeton, N.J.: Princeton University Press.
- Wegner, D.M., 1995. A computer network model of human transactive memory. *Social Cognition*, 13(3): 319-339.
- Yuan, Y.C. 1990. *Multiple Imputation for Missing Data: Concepts and New Developments*. Paper 267-65. Rockville, MD: SAS Institute Inc.