

## Destabilizing Networks

Kathleen M. Carley  
Ju-Sung Lee  
David Krackhardt  
Carnegie Mellon University  
November 2001

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

This 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.

## Destabilizing Networks

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, modeling 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.

#### [1] 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). These 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.
2. An individual or group that is unlikely to act even if given alternative information.
3. An individual or group that if given new information can propagate it rapidly.
4. An individual or group that has relatively more power and can be a possible source of trouble, potential dissidents, or potential innovators.
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).

6. An individual, group, or resource that provides redundancy in the network.

## [2] 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 these approaches is that they search for behavior that is different from some baseline such as random networks, biased networks, or a sample of existing networks. For example, the following kinds of patterns or breaks in patterns can be examined:

1. The basic components that account for the networks structure; e.g., the number and types of sub-groups, or the number of triads, stars, and the extent of reciprocity.
2. The central tendency within a set of networks, and the networks that are anomalous when contrasted with the other networks in the set.
3. Critical differences between two or more sets of networks; e.g., are programming teams structured differently than sales teams or are managers' mental models different from subordinates.
4. Which components in the network are structured significantly differently from the rest of the overall network.
5. Whether the existing network is coherent; i.e., what is the likelihood that there are key missing nodes or relations.

## [3] 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. 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 are a series of illustrative examples of potential applications where various researchers in this area have worked or are working:

1. Suggesting factors that make groups adaptive or maladaptive.
2. Examining the efficacy of different policies for destabilizing networks; e.g., what kinds of networks can be destabilized by simply removing the leader? What are the characteristics of networks that are difficult to destabilize?
3. 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?

4. Predicting the rate of information diffusion and the impact of different technologies for spreading information and so changing beliefs through social influence processes.
5. Predicting voting outcomes or likelihood of consensus in groups, given the existing social networks and initial beliefs.
6. 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.
7. Determining how close your group or company is to having its core competencies and processes discovered by another group (i.e. inevitable disclosure).
8. Examine the efficacy of different marketing and information warfare strategies.

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 technologies 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 often limits analysis to small networks, those less than a few hundred nodes. Second, we have no public databases of large networks on which to test new technologies. 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. 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 networks. However, these are not de-coupled systems. Humans learn when they interact with each other and what they learn changes with whom they interact. 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.

We have built a relatively simple computational model of this dynamic process – CONSTRUCT-O (see for a description of this model, Carley & 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.

These changes also alter whether or not there is an emergent leader and which agent takes on that role. An emergent leader is 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 & Ren, 2001). Individuals with high cognitive load 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 who they know. 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.

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.

Using this model we examine two very distinct structures – a hierarchical centralized structure and a distributed decentralized one. The Krackplot representations of these stylized 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. 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). The “Emergent Leader” agent is denoted by a rectangular node labeled ‘LEADER.’ This agent is the individual with the highest cognitive load (i.e. most resources, tasks, and communication/network ties). The agent with the most network ties is denoted by an oval node labeled ‘CENTRAL.’ If the agent is both the emergent leader and the most central then that agent is denoted by a rectangular node labeled ‘LEADER/CENTRAL.’ 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.

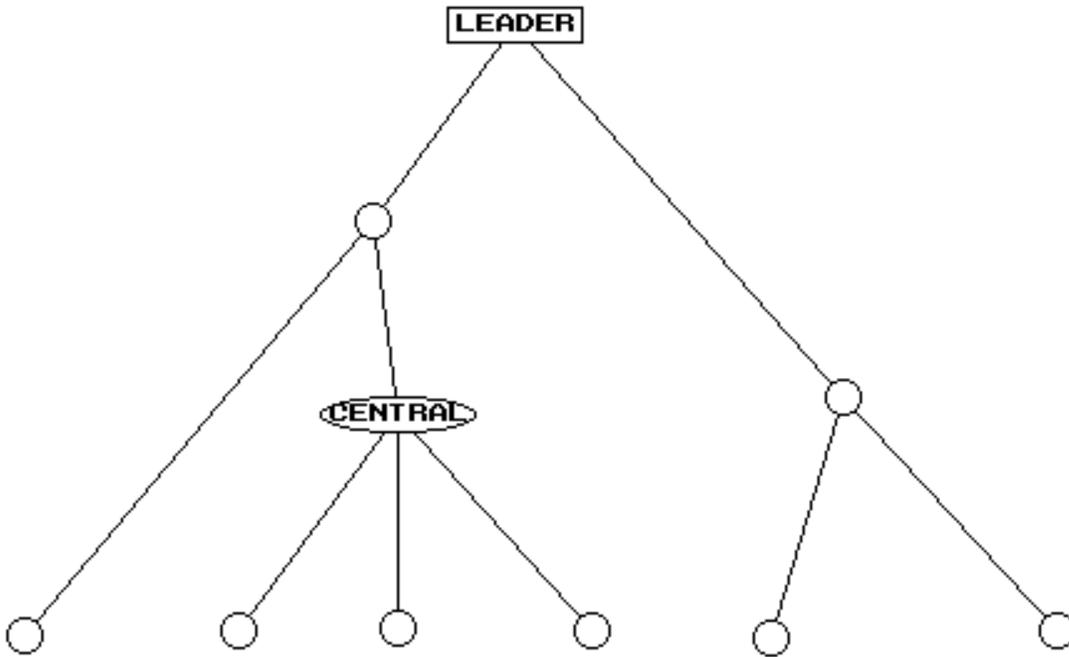


Figure 1. A Stylized Hierarchical Centralized Network

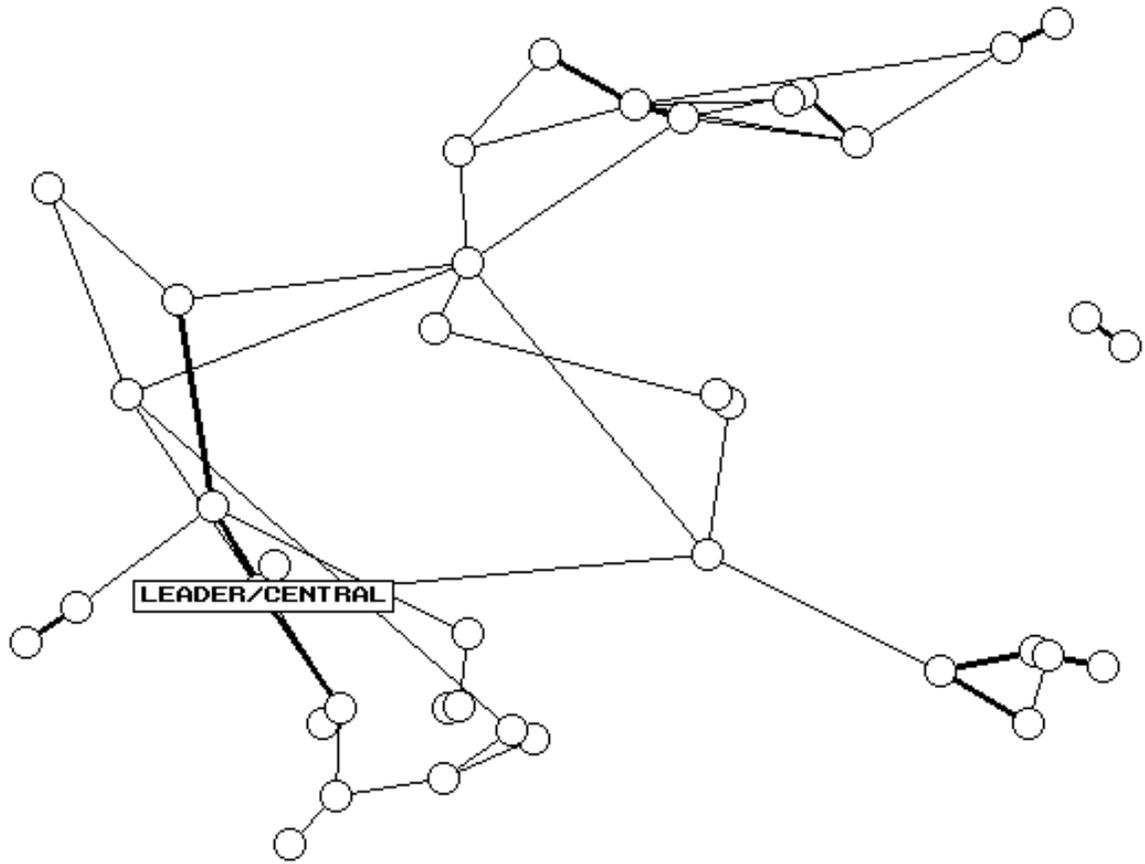


Figure 2. A Stylized Distributed Decentralized Network

Clearly, for these idealized structures, it would be relatively easy to destabilize them as compared to a large human network. Nevertheless, even this highly stylized example illustrates that the problem of destabilization is more difficult for a distributed than for a hierarchical network. Most people when they visually examine the hierarchical network will predict that removal of the LEADER is likely to destabilize the network. In the hierarchal network (Figure 1) the emergent leader is not necessarily the most central agent. When the LEADER communicates with one of the subordinates that increases his/her centrality. Thus another candidate agent whose removal might destabilize the hierarchical network is the CENTRAL agent. In the particular hierarchal network in Figure 1 it is the removal of the LEADER that destabilized the network. In fact, the hierarchal network, Figure 1, can be easily destabilized by removing, isolating, hiring away the LEADER or by stopping the flow of information or resources through all links connected to the LEADER. For the hierarchal network, the LEADER's ability to control the hierarchy can also be decreased by adding new links.

However, 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 possible to destabilize the network. This is the case even when, as in Figure 2, the emergent leader is the most central agent. 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.

Now consider the case of adaptation. Since individuals can learn, the underlying social networks are dynamic. When nodes are removed or isolated the structure of the network changes in response. As a result, removing a node may result in a new individual emerging as the leader. The possible path of change for the hierarchal network is shown in Figure 3 and for the distributed network 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.' In Figure 3a the initial hierarchical network is shown. Then, the emergent leader, is extracted. Figure 3b contains the resultant network that emerges after the original emergent leader is removed. Immediately, the extraction of the agent – LEADER - causes the hierarchy to break up in to two smaller networks. Once the LEADER is extracted the network reforms with two emergent leaders who are essentially competing for control. In Figure 3c the network has adapted to the loss and a new single leader has emerged. Not all hierarchies will change in this way – but this diagram is illustrative of the impact of extracting a leader on a hierarchical network.

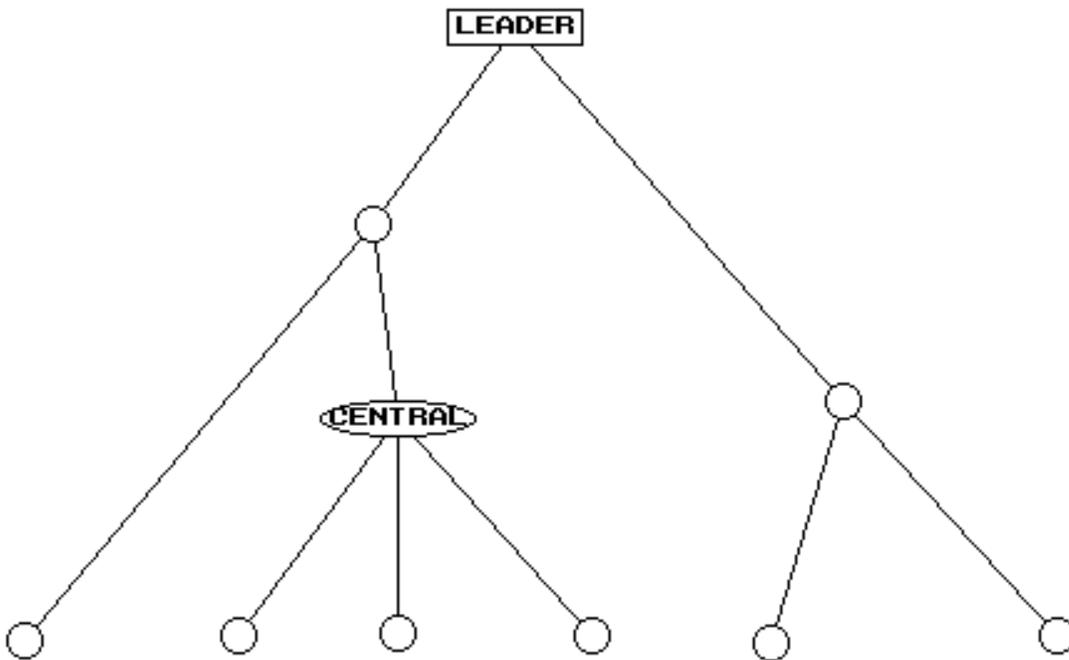


Figure 3a. Removal of an Emergent Leader in a Stylized Hierarchical Centralized Network – Initial Structure

REMOVED

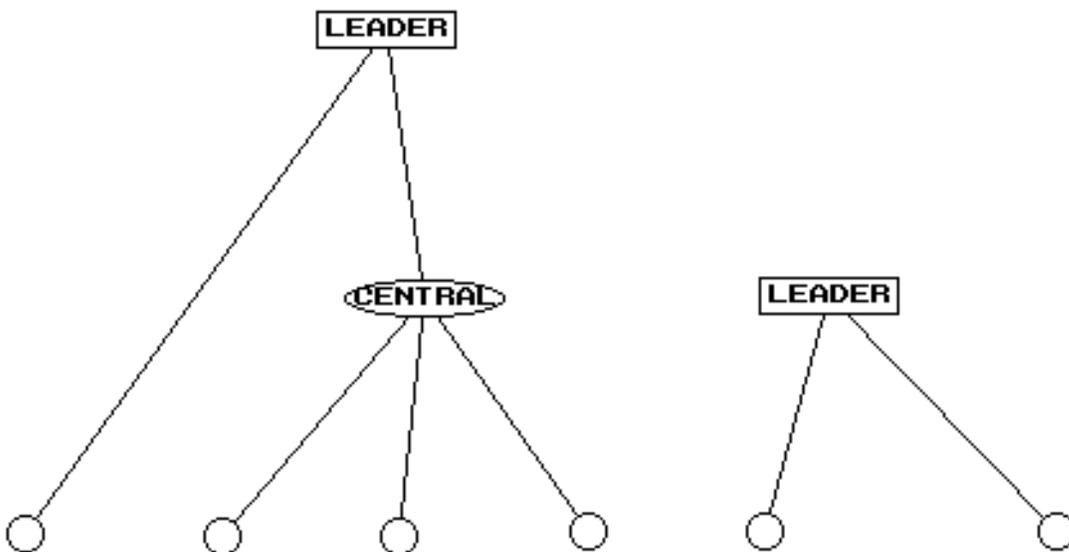


Figure 3b. Removal of an Emergent Leader in a Stylized Hierarchical Centralized Network – Immediate Response to Removal of Emergent Leader

REMOVED

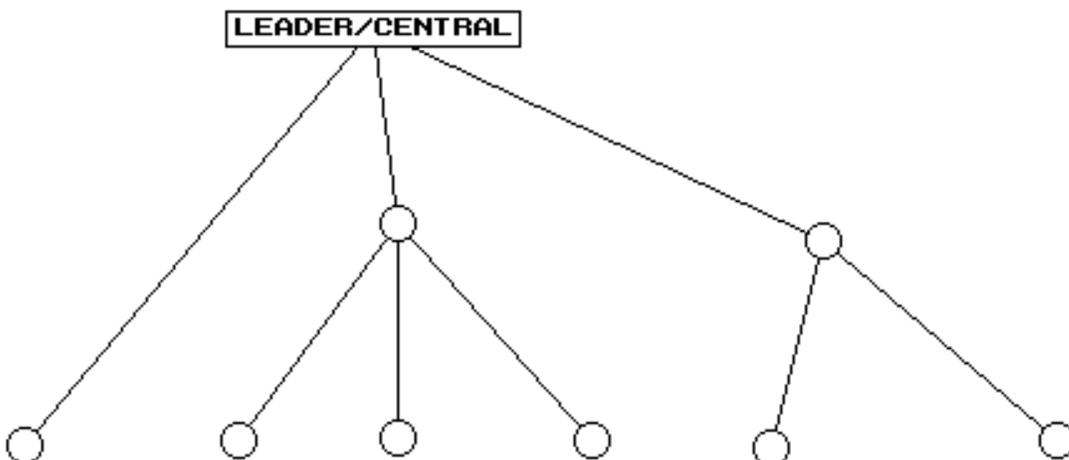


Figure 3c. Removal of an Emergent Leader in a Stylized Hierarchical Centralized Network – Eventual Response

In Figure 4, the consequences of removing an emergent leader on a distributed decentralized network are portrayed. The initial structure is in Figure 4a. As in Figure 3b, the leader, LEADER/CENTRAL, is now extracted. In Figure 4b, like Figure 3b, the position that the original leader would have held if he/she had not been extracted is denoted by the word removed. In Figure 4b we see that after that a new leader emerges in the same vicinity as the 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 4c, multiple new leaders emerge. In addition, the agent who in Figure 4b 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 suggests that the structure of the network in that vicinity promotes the development of emergent leaders. Further, when the original leader was present, that individual was inhibiting the emergence of alternative leaders. The original leader had maintained key resources 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, multiple leaders could eventually emerge.

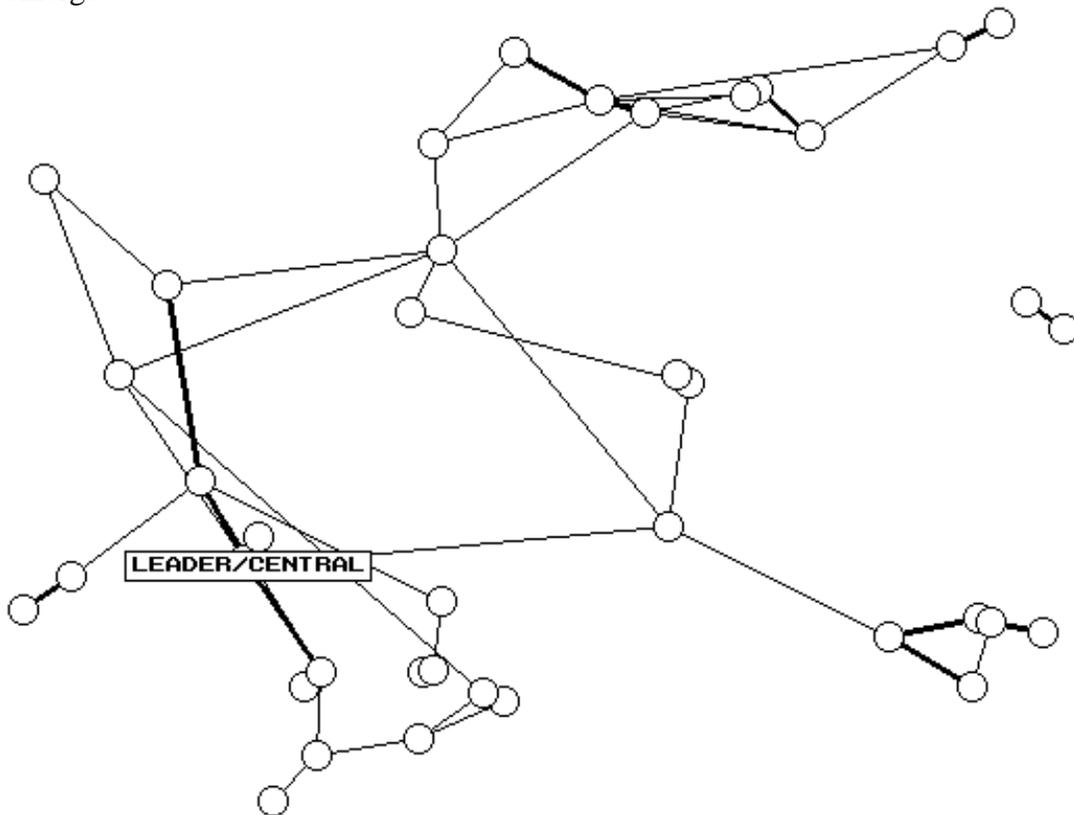


Figure 4a. Removal of an Emergent Leader in a Stylized Distributed Decentralized Network – Initial Structure

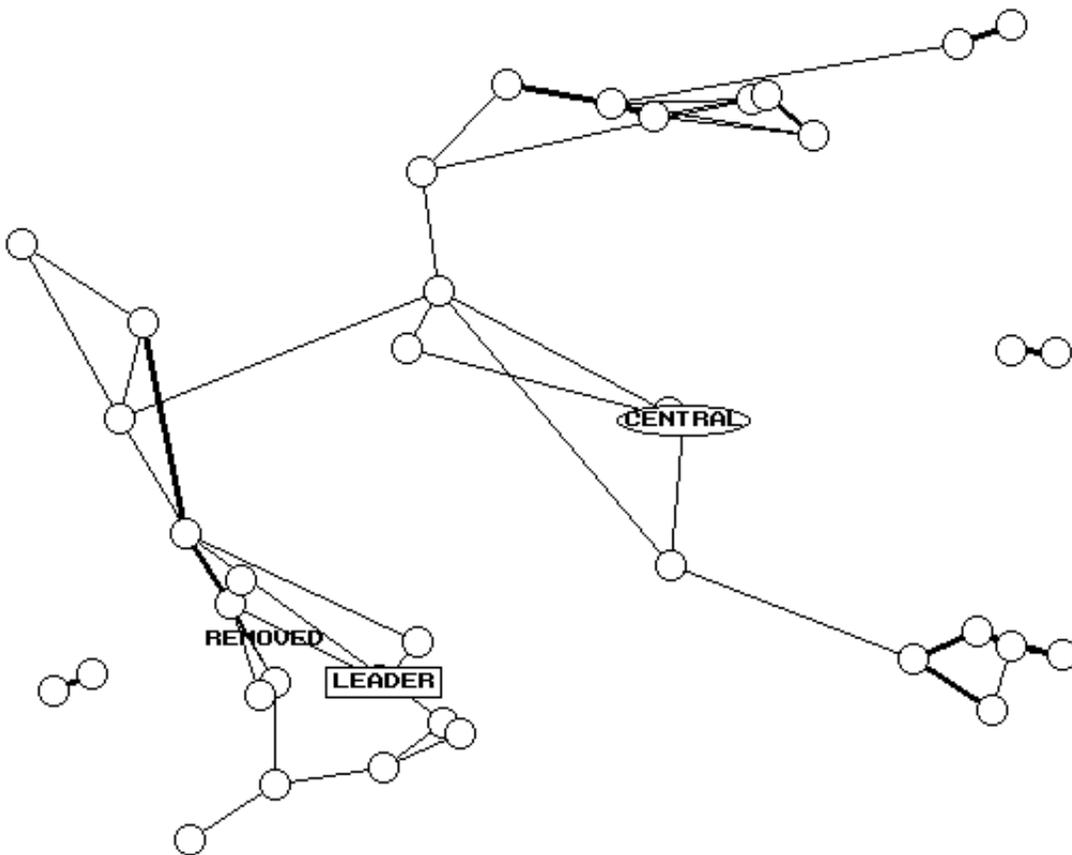


Figure 4b. Removal of an Emergent Leader in a Stylized Distributed Decentralized Network – Immediate Response to Removal of Emergent Leader

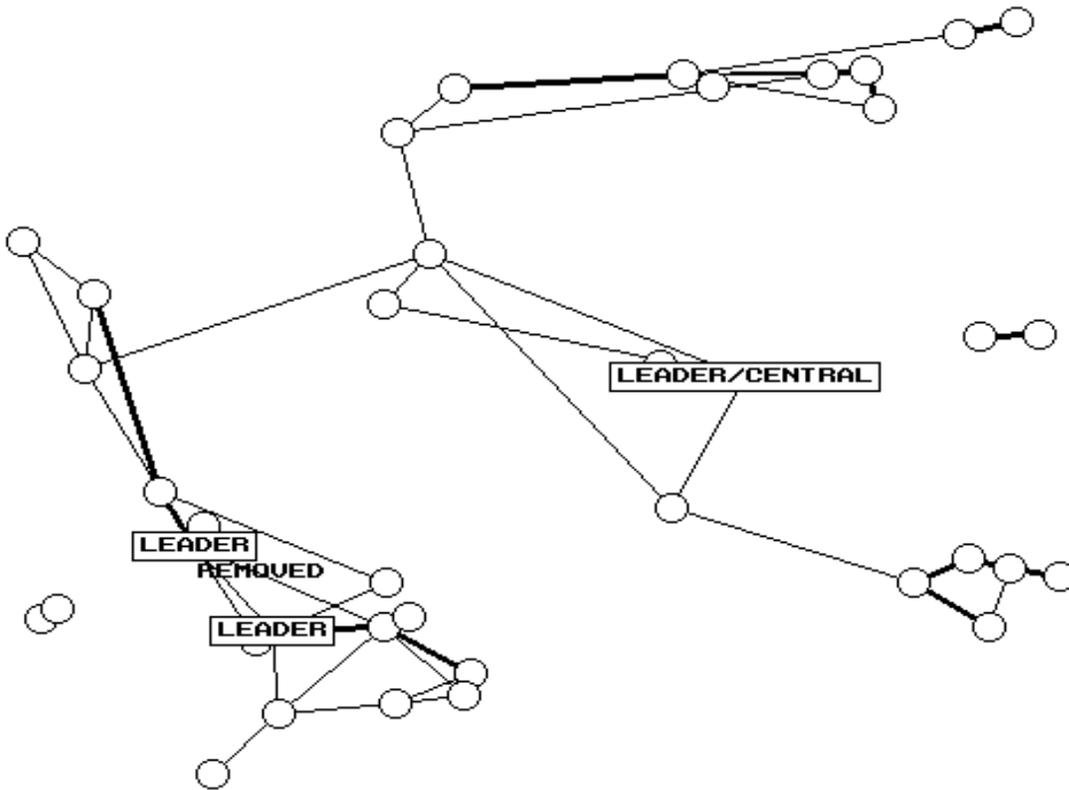


Figure 4c. Removal of an Emergent Leader in a Stylized Distributed Decentralized Network – Eventual Response

As Figure 4 suggests, it may be very difficult to destabilize a decentralized structure. As one leader is isolated another rises to take his or her place. The dynamics are 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 & Hill, 2001, Krackhardt & Carley, 1998). 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.

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. 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. When the position of the agent's in the meta-network is considered, it becomes clear that it is the overall position in the meta-network, and not just in the social network subcomponent that is key. 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. 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). Overall cognitive load, not simply structural power, is key to tracking who the emergent leader is likely to be.

Table 1. Simplified Meta-Matrix Representation of the Meta-Network			
	Agents	Knowledge	Tasks
Agents	Social Network	Knowledge Network	Assignment Network
Knowledge		Information Network	Needs Network
Tasks			Task-Precedence Network

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 incredibly 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-colored networks with high levels of missing data. The development of such tools are needed if we are to successfully meet the challenge of understanding, predicting and explaining the behavior 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

- Carley, Kathleen M. 2001. “Computational Organization Science: A New Frontier.” Paper prepared for the NAS sponsored Sackler Colloquium, Irvine CA, October 2001.
- Carley, Kathleen M. 2000. “Information Security: The Human Perspective.” CASOS Working Paper, Carnegie Mellon University
- Carley, Kathleen 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, Kathleen, 1990, “Group Stability: A Socio-Cognitive Approach.” Pp. 1-44 in Lawler E., Markovsky B., Ridgeway C. & Walker H. (Eds.) *Advances in Group Processes: Theory & Research* . Vol. VII. Greenwich, CN: JAI Press.
- Carley, Kathleen, 1991. “A Theory of Group Stability.” *American Sociological Review* , 56(3): 331-354.
- Carley, Kathleen M. & Yuqing Ren, 2001, “Tradeoffs Between Performance and Adaptability for C<sup>3</sup>I Architectures.” In *Proceedings of the 2001 Command and Control Research and Technology Symposium*, Annapolis, Maryland, June, 2001.
- Carley, Kathleen M. & Vanessa Hill, 2001, “Structural Change and Learning Within Organizations”. In *Dynamics of Organizations: Computational Modeling and*

- Organizational Theories. Edited by Alessandro Lomi and Erik R. Larsen, MIT Press/AAAI Press/Live Oak, Ch. 2. pp 63-92.
- Carley, Kathleen M. & Les Gasser, 1999, "Computational Organization Theory." Ch. 7 in Gerhard Weiss (Ed.) *Distributed Artificial Intelligence*. Cambridge, MA: MIT Press.
- Epstein, Josh & Rob Axtell, 1997. *Growing Artificial Societies*, Boston, MA: MIT Press.
- Krackhardt D. 1990. Assessing the Political Landscape: Structure, Cognition, & Power in Organizations. *Administrative Science Quarterly*, 92, 142-154.
- Krackhardt, David & Kathleen M. Carley, 1998, "A PCANS Model of Structure in Organization" Pp. 113-119 in *Proceedings of the 1998 International Symposium on Command and Control Research and Technology*. June. Monterey, CA.