

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.

