Dynamic Network Analysis

Kathleen M. Carley
Institute for Software Research International
Carnegie Mellon University

Abstract

Dynamic network analysis (DNA) varies from traditional social network analysis in that it can handle large dynamic multi-mode, multi-link networks with varying levels of uncertainty. DNA, like quantum mechanics, would be a theory in which relations are probabilistic, the measurement of a node changes its properties, movement in one part of the system propagates through the system, and so on. However, unlike quantum mechanics, the nodes in the DNA, the atoms, can learn. An approach to DNA is described that builds DNA theory through the combined use of multi-agent modeling, machine learning, and meta-matrix approach to network representation. A set of candidate metric for describing the DNA are defined. Then, a model built using this approach is presented. Results concerning the evolution and destabilization of networks are described.

Acknowledgement

The research reported herein was supported by the National Science Foundation NSF IRI9633 662, the Office of Naval Research (ONR) Grant No. N00014-97-1-0037 and Grant No. 9620.1.1140071, Additional support was provided by the NSF IGERT 9972762 for research and training in CASOS and by 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 authors 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.

Citation

Kathleen M. Carley, forthcoming, “Dynamic Network Analysis” in the Summary of the NRC workshop on Social Network Modeling and Analysis, Ron Breiger and Kathleen M. Carley (Eds.), National Research Council.

Dynamic Network Analysis

Terrorist organizations have network structures that are distinct from those in typical hierarchical organizations – they are cellular and distributed. While most commanders, politicians and intelligence agents have at least an intuitive understanding of hierarchies and how to affect their behavior, they have less of an understanding of how to even go about reasoning about dynamic networked organizations (Ronfelt and Arquilla, 2001). It is even more difficult for us to understand how such networks will evolve, change, adapt and how they can be destabilized.

Clearly social network analysis can be applied to the study of covert networks (Sparrow, 1991). Many are stepping forward suggesting that to understand these networks we just need to "connect the dots" and then isolate the "key actors who are often defined in terms of their "centrality" in the network. To an extent, this is right. However, it belies the difficulty of "connecting the dots" in terms of mining vast quantities of information, pattern matching on agent characteristics for people who go under multiple aliases, and still ending up with information the may be intentionally misleading, inaccurate, out-of-date, and incomplete. Further, this belies the difficulty in "knowing" who is the most central when you have at best only a sample of the network. Finally, and critically, this approach does not contend with the most pressing problem – the underlying network is dynamic. Just because you isolate a key actor today does not mean that the network will be destabilized and unable to respond. Rather, it is possible, that isolating such an actor may have the same effect as cutting off the Hydra's head; many new key actors may emerge (Carley, Lee and Krackhardt, 2001).

To understand the dynamics of terrorist, and indeed any, network we need to understand the basic processes by which networks evolve. Moreover, we have to evaluate isolation strategies in the face of an evolving network and in the face of missing information. To ignore either the dynamics or the lack of information is liable to lead to erroneous, and possibly devastatingly wrong, policies. Taking in to account both the dynamics and the lack of information should engender a more cautious approach in which we can ask, "if we do x what is likely to happen?"

Limitations to Traditional SNA

Traditionally, social network analysis (SNA) has focused on small, bounded networks, with 2-3 types of links (such as friendship and advice) among one type of node (such as people), at one point in time, with close to perfect information. To be sure there are a few studies that have considered extremely large networks, or two types of nodes (people and events), or unbounded networks (such as inter-organizational response teams); however, these are the exception not the norm. However, such studies are still the exception not the rule. Further, while it is understood, at least in principle how to think about multi-modal, multi-plex, dynamic networks, the number of tools, the interpretation of the measures, and the illustrative studies using such "higher order" networks are still in their infancy relative to what is available for simpler networks. Finally, many of the tools do not scale well with the size of the network or degrade gracefully with errors in the network; e.g., they may be too computationally expensive or too sensitive to both type 1 and 2 errors. What is needed is a dynamic network analysis theory and toolkit. We are working to develop such a tool kit and the associated metrics and decision aids. In this paper, one such tool, DyNet is described and used to examine various isolation strategies.

Dynamic Network Analysis

Recently there have been a number of advances that extend SNA to the realm of dynamic analysis and multi-color networks. There are three key advances: 1) the meta-matrix, 2) treating ties as probabilistic, and 3) combining social networks with cognitive science and multi-agent systems. These advances result in a dynamic network analysis.

Meta-Matrix: Carley (2002) combined knowledge management, operations research and social networks techniques together to create the notion of the meta-matrix – a multi-color, multiplex representation of the entities and the connections among them. The Meta-matrix is an extension and generalization of the PCANS approach forwarded by Carley and Krackhardt (1999) that focused on people, resources and tasks. For our purpose, the entities of interest are people, knowledge/resources, events/tasks and organizations – see table 1. This defines a set of 10 inter-linked networks such that changes in one network cascade into changes in the others; relationships in one network imply relationships in another. For example, co-membership in an organization or co-attendance at an event for two people suggests a tie in the social network between these two people. A group, such as a terrorist network, can be represented in terms of an overtime sequence of such networks. In fact, any organization or group can be represented in this fashion and we have used this representation on numerous occasions to characterize actual organizations and to predict their ability to adapt.

All graph theory and network measures can be defined in terms of whether they can or have been applied to which cells. Further, on the basis of this meta-matrix new metrics can be developed that better capture the overall importance of an individual, task, or resource in the group. An example of such a metric is cognitive load – the effort an individual has to employ to hold his role in the terrorist group - and it takes in to account, who he interacts with, which events he has been at, which organizations he is a member of, the coordination costs of working with others in the same organization or at the same event or in learning from an earlier event or training for an upcoming event. A large number of such metrics have been developed and analyzed in terms of their ability to explain the evolution, performance, and adaptability of dynamic networks.

A key difficulty from a growth of science perspective, is that as we move from SNA to DNA the number, type, complexity, and value of measures changes. A core issue for DNA is what are the appropriate metrics for describing and contrasting dynamic networks. Significant new research is needed in this regard. To date, our work suggests that a great deal of leverage can be gained in describing networks by focusing on measures that utilize more of the cells in the meta-matrix. For example, cognitive load, which measures the cognitive effort and individual has to do at one point in time has been shown to be a valuable predictor of emergent leadership (Carley and Ren, 2001). Cognitive load is a complex measure that takes into account the number of others, resources, tasks the agent needs to manage and the communication needed to engage in such activity. In addition, we find that for any of the cells in the meta-matrix, particularly for large scale networks, many of the standard graph level measures have little information content as the network grows in size (Anderson, Butts and Carley, 1999) and/or are highly correlated with each other. A set of measures that are generally not correlated, scale well, and are key in characterizing a network are the size of the network (number of nodes), density (either as number of ties or the typical social network form number of ties/number of possible ties), homogeneity in the distribution of ties (e.g., the number of clusters or subcomponents, the variance in centrality), rate of change in nodes, and rate of change in ties. The point is not that these are the only measures needed to characterize dynamic networks. The point is that these are a candidate set that have value and that as a field we need to develop a small set of metrics that can be applied to networks, regardless of size, to characterize the dynamics.

Table 1. Meta-Matrix

	People	Knowledge/Resources	Events/Tasks	Organizations
People	Social network	Knowledge network	Attendance network	Membership network
Knowledge/Resources		Information network	Needs network	Organizational capability
Events/Tasks			Temporal ordering	Institutional support or attack
Organizations				Inter-organizational network

Probabilistic Ties: The ties in the meta-matrix are probabilistic. Various factors affect the probability, including the observer's certainty in the tie and the likelihood that the tie is manifest at that time. Bayesian updating techniques (Dombroski and Carley, 2002), cognitive inferencing techniques, and models of social and cognitive change processes (Carley, 2002; Carley, Lee and Krackhardt, 2001) can be used to estimate the probability and how it changes over time. We are in the process of exploring techniques for combining the cognitive inferencing with the cognitive change process models.

Multi-Agent Network Models: A major problem with traditional SNA is that the people in the networks are not treated as active adaptive agents capable of taking action, learning, and altering their networks. There are several basic, well known, social and cognitive processes that influence who is likely to interact with whom: relative similarity, relative expertise, and co-worker. Carley uses multi-agent technology in which the agents use these mechanisms, learn, take part in events, do tasks to model organizational and social change. The dynamic social network emerges from these actions. The set of networks linking people, knowledge, tasks and other groups or organizations co-evolve. Carley, Lee and Krackhardt (2001) use simple learning mechanisms to dynamically adjust networks as the agents in them attended events, learned new information, or were removed from the network. In DyNet, described herein, additional mechanisms center on agent isolation are also considered.

DNA has a wide range of applications. For example, this approach is being used to examine the likely impact of unanticipated events in the VISTA project (Diedrich et al, forthcoming), the possible effects of biological attacks on cities in BioWar (Carley et al, 2002), in evaluating CIO response strategies to denial of service attacks (Chen, 2002), and evaluating information security within organizations – ThreatFinder Project (Carley, 2001). See also www.casos.ece.cmu.edu current projects and working papers.

Dynamic Network Theory

To move beyond representation and method, we need to ask, "How do networks change?" What are the basic processes? From the meta-matrix perspective, the processes are easy – things that lead to the adding and dropping of nodes and/or relations – see table 2. Again, no claim is being made that the processes listed in table 2 cover the complete spectrum; rather, they illustrate the types of node change processes that need to be postulated. A full theory of dynamic networks needs to speak to such mechanisms.

Table 2. Basic Change Processes for Nodes in the Meta-Matrix

People	Knowledge/Resources	Events/Tasks	Organizations
Birth Death Promotion Mobility Recruitment Incarceration Isolation	Innovation Discovery Forgetting Consumption	Goal Change Re-engineering Development of new technology Stop usage of technology	Organizational birth Organizational death Mergers Acquisitions Legislation of new entity

Similarly, there are a set of processes that lead to the addition and removal of relations. Basic processes are cognitive, social and political in nature. Cognitive processes have to do with learning and forgetting, the changes that occur in ties due to changes in what individuals know. Social changes occur when one agent or organization dictates a change in ties, such as when a manager re-assigns individuals to tasks. Finally, political changes are due to legislation that effect organizations and the over-arching goals. To illustrate what is meant, a limited number of such processes are described in Table 3. Further, and this should be obvious, processes that add or eliminate nodes also affect relations to/from that node. For example, if all individuals in a society forget a particular piece of information that knowledge node, no longer exists and all connections from people to it are now eliminated.

Table 3. Change Processes for Relations in the Meta-Matrix

	People	Knowledge/ Resources	Events/ Tasks	Organizations
People	Motivation to Interact Change in access	Learning Acquisition	Re-assignment	Mobility Recruitment
Knowledge/ Resources		Discovery Analogical reasoning	Innovation	IP development
Events/Tasks				Re-engineering Out-sourcing
Organizations				Alliances Coalitions

DyNet

The purpose of the DyNet project is to develop the equivalent of a flight simulator for reasoning about dynamic networked organizations. Through a unique blending of computer science, social networks and organization theory we are creating a new class of tools for managing organizational dynamics. The core tool is DyNet – a reasoning support tool for reasoning under varying levels of uncertainty about dynamic networked and cellular organizations, their vulnerabilities, and their ability to reconstitute themselves. Using DyNet the analyst would be able to see how the networked organization was likely to evolve if left alone, how its performance could be affected by various information warfare and isolation strategies, and how robust these strategies are in the face of varying levels of information assurance.

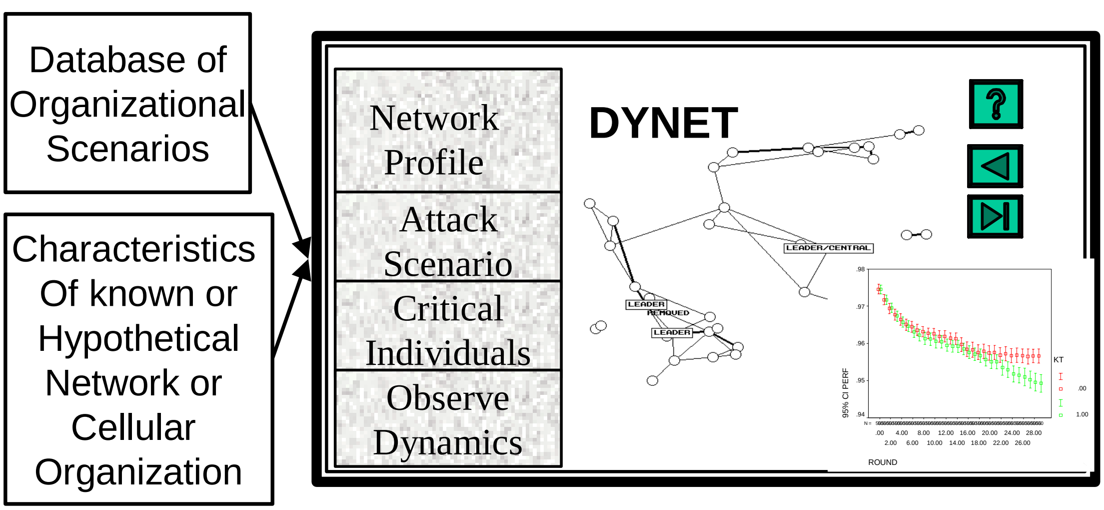

Figure 1. DYNET: A desktop tool for reasoning about dynamic networked and cellular organizations.

DyNet is intended to be a desktop system that can be placed in the hands of intelligence personnel, researchers, or military strategists. Through hands-on what if analysis the analysts will be able to reason in a what –if fashion about how to build stable adaptive networks with high performance and how to destabilize networks. There are many applications for such a tool including: threat assessment; assessing information security risks in corporations; intel training; simulation of the red team in a gaming situation, and estimation of efficacy of destabilization policies. Currently an alpha version exists as a batch program (no visualization) and it has been used to evaluate simple isolation strategies. The system can handle data on real networks.

The DyNet tool is a step toward understanding how networks will evolve, change, adapt and how they can be destabilized. The goal will be to incorporate all of the evolutionary mechanisms previously discussed. DyNet, which is a computer model of dynamic networks, can also be thought of as the embodiment of a theory of dynamic networks. The focus of this theory is on the cognitive, and to a lesser extent, social processes by which the networks in the meta-matrix evolve. The basic cognitive forces for change in DyNet are learning, forgetting, goal-setting, and motivation for interaction. The basic social forces for change are recruitment, isolation, and to a limited extent the initiation of rumors and training.

The basic motivations for interaction are relative similarity, relative expertise or some combination of the two. Relative similarity is based on the fundamental finding of homophilly, the tendency of interacting partners to be similar. Arguments surrounding this fundamental process include the need for communicative ease, comfort, access, and training. Relative expertise is based on the fundamental finding that when in doubt people will turn they view as experts for information. Arguments surround this fundamental processes include the need to acquire, desire to minimize search, desire to optimize information, and so on. Other basic motivations such as the need to exhibit competence and the need to coordinate have also been identified and will be added to DyNet but are not in the current system.

Among the attrition strategies are removal of the most "central" individual, removal of the individual with the highest cognitive load, and removal of individual's at random. User's can control the frequency and severity of such attrition strategies. Previous studies using this system have shown that a) it is difficult to completely destabilize a network, b) that the best strategy depends on the structure of the network, and c) attrition strategies vary in whether there effectiveness is enhanced or diminished by removing multiple agents at once or sequentially (Carley, 2002).

Agents can be distinguished based on fixed characteristics such as race, family and gender, and on knowledge (or training). Further, the agents can operate in a world without information technology or augmented by access to email, web pages, or manuals. Access to others can be restricted, as might be the case when operatives live in different countries. Performance metrics include task completion, accuracy, energy for tasks, information diffusion, and group cohesion. Finally, the basic networks can be extracted continually in order to see the system evolve. Among the networks that can be extracted are the knowledge network, the overall social network, the emotive or "friendship" networks, and the acquisition or "advice" network. The network evolutionary strategies include learning (during interaction), forgetting, personnel attrition, misinformation, and changing task demands. DyNet offers the user the choice of entering specific networks or entering network characteristics (such as size and density).

Results

Using DyNet a series of virtual experiments were run. These experiments were designed to examine the interaction between network structure, dynamics (particularly in response to isolation), and the information that the observer has on which to base the isolation strategies. In figure 2, a very high level conceptualization of these differences is shown. Three possible isolation strategies: isolating individuals at random, isolating those who are the most central (degree centrality), and isolating those with the highest cognitive load are shown relative to a specific organization and networks within it. Given that the networks are evolving at issue is which of these strategies will be the most effective? Further, we might ask, if the social network was different, e.g., less hierarchical, would that matter?

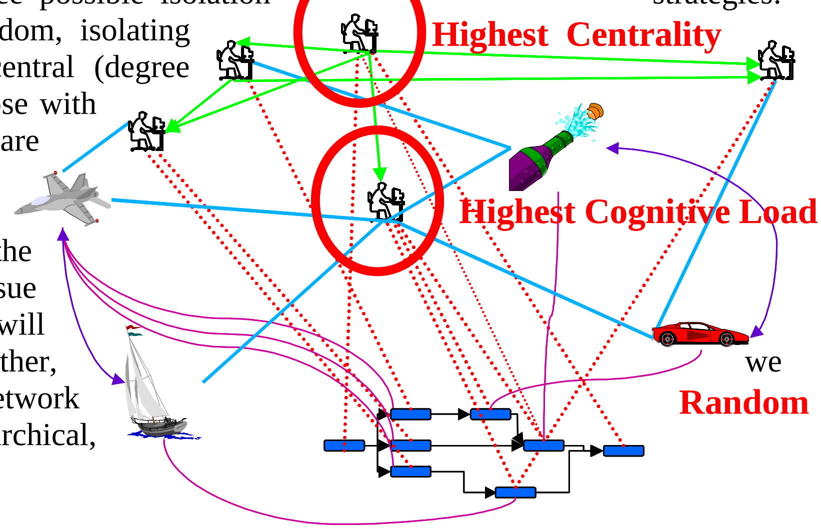

Figure 2. Structure, Isolation and Dynamics

The Structure of the Network Matters

The first finding, and it is quite robust, is that the structure of the network matters. That is, random networks in which the relations are distributed in an independent and identical fashion, hierarchies, and cellular networks all evolve quite differently, require different strategies to destabilize, have different abilities to diffuse information, and exhibit different performance for the same task. In Figures 3 and 4, this difference is illustrated with respect to the networks ability to recover from isolation strategies. In figure 3 we see the impact of the three isolation strategies on a random iid network and in figure 4 the impact of the same strategies on a cellular network. As an aside, the particular cellular network simulated here is one whose features map onto available information about covert networks, such as the cells are completely connected internally and cell size ranges from 3-10 members. In these figures not only do we see that the isolation strategies vary in their effectiveness based on the structure of the network they are attacking, but in addition, cellular networks are able to recover from the attacks.

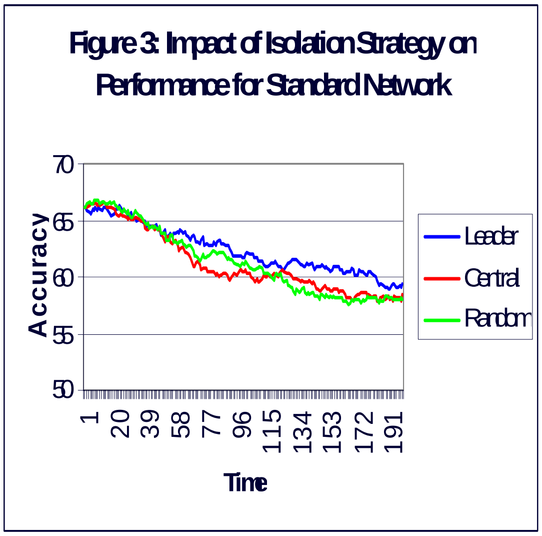

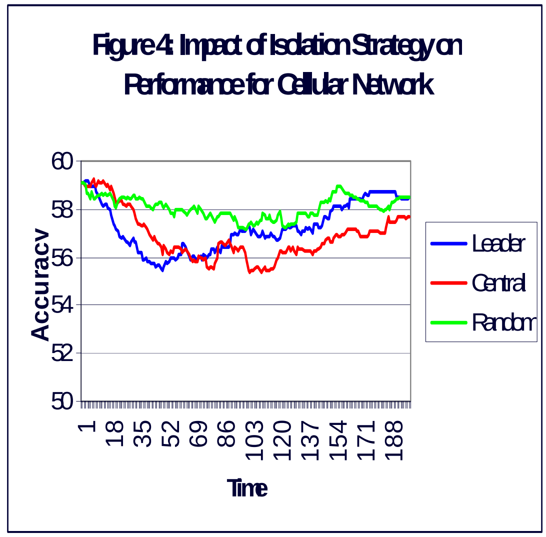

Networks Can Heal Themselves

A second key finding is that networks are generally able to heal themselves. That is isolation of a node that links disparate groups together typically does not leave those groups disconnected. Rather the basic social and cognitive processes outlined lead individuals to seek alternative contact points to interact with. For example in Figure 5 we see on the left a network where the person with the highest cognitive load, the emergent leader was isolated. A consequence is that multiple new leaders emerge, each of whom ends up being more directive than the original leader. Healing is not guaranteed and in fact depends on the underlying structure, the cultural basis for interaction, the degree of isolation, the frequency of isolation, and the strategy for isolation. For example, as was seen in Figure 4, cellular networks heal themselves regardless of which isolation strategy is used against it. In this case, the cell structure of the network enables the network as a to engage in what appears as "meta-learning," i.e., learning how to recover from unanticipated attrition. Cellular networks, which are the structure most like those used by terrorist organizations, are very difficult to destabilize. The reasons are complex, but a key factor is that such network structures are able to heal relatively faster than other structure both in terms of the re-emergence of leaders and in terms of performance recoveries after personnel have been removed.

Full Information is Not Necessary

In the foregoing two examples, we saw the impact of destabilization strategies on network without considering "how is it that we know what we know?" Or in other words, "if we are not sure what the underlying network looks like, how confident can we be in our predictions about how to destabilize it?" Notice, that in traditional SNA, typically we have close to full information. For covert networks we do not. Information may be missing because we don't know some of the nodes – the people involved, or because we don't know some of the relations.

Figure 5. The network before (left) and after (right) the isolation of the leader.

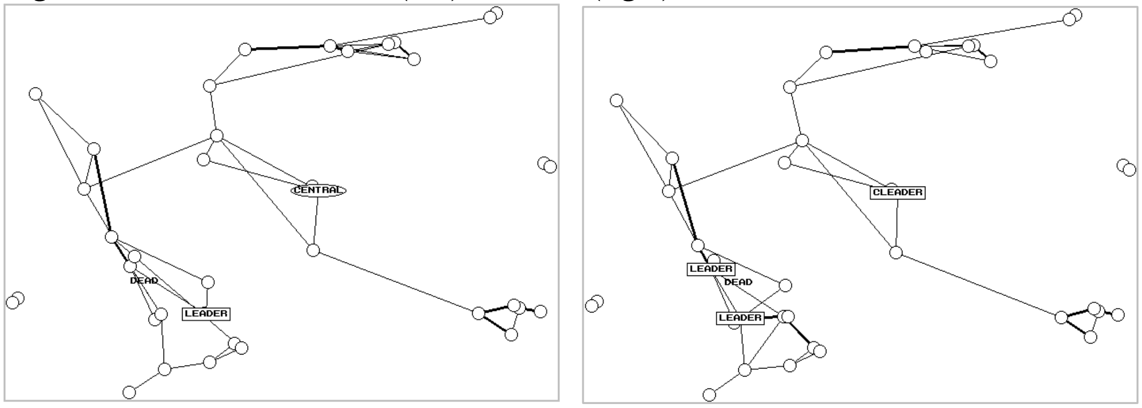

Figure 6. The impact of incomplete information about who is in the network.

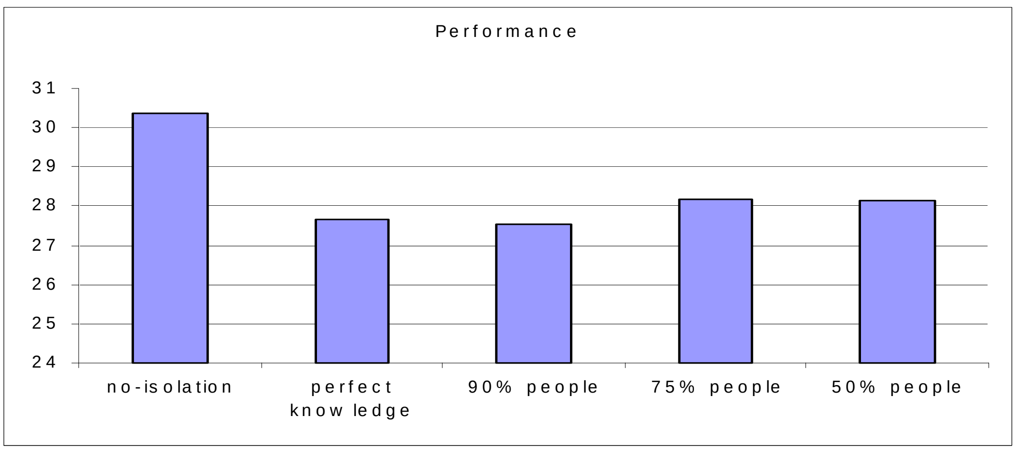

In Figure 6 we see the impact of not knowing al the nodes. Here we see the comparison of no attack, versus the average impact of isolating personnel (across the three isolation strategies) under conditions of full information, knowing 90% of the nodes, 75% and only 50%. Clearly having close to perfect or perfect knowledge means that more effective isolation strategies are found. Note, however, that any isolation is better than none, assuming our goal is to degrade the performance and that we don't need perfect information to be quite effective.

Now, consider the case where we don't have perfect information about the relations. One way for this to occur is if we don't know all the knowledge or resources that are available to the network. In figure 7 we see the impact of having imperfect knowledge of the relations as a function of how much do we know about the other entities, in this case what there is for the other to know. Here we see that we actually do better knowing less. This is due to the interaction between what we know and the isolation strategy. Essentially, when we don't really know the underlying social and knowledge network we may overestimate the primacy of a person, who although not the key in terms of degree centrality, is more central in terms of cognitive load. Thus, in effect, less knowledge makes both the centrality and the cognitive load strategies more similar resulting in on average lower performance due to the fact that cellular networks are more devastated by the extraction of such emergent leaders, at least in the short run. Further, reduced information about relations makes all isolation strategies more mixed thus inhibiting the ability of the opponent to engage in meta-learning.

Figure 7. The impact of incomplete information about what people know.

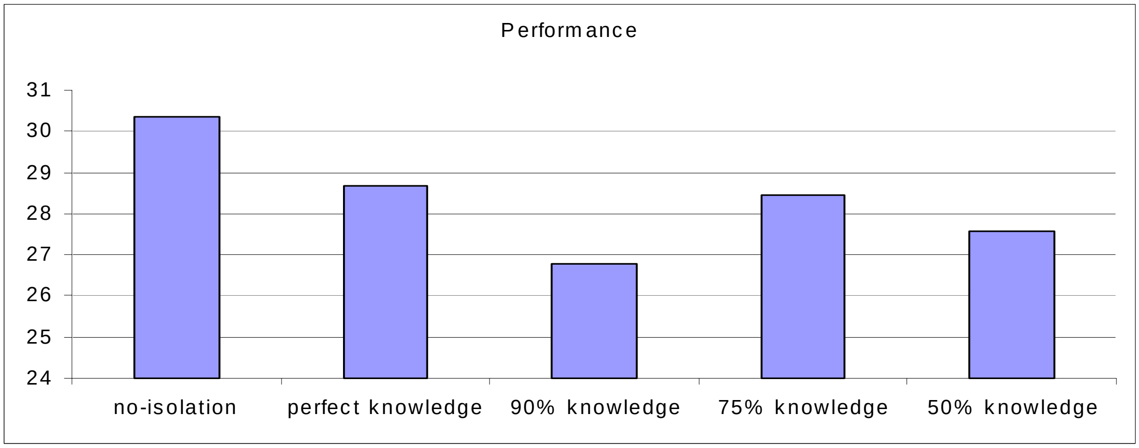

Summary

Thinking about networks from a dynamic perspective is absolutely essential to understanding the modern world. An approach toward dynamic networks has been outlined. There are several distinctive hallmarks to this approach. First, in contrast to other multi-agent work, the agents we describe are in actual social networks. Here, the networks and the agents co-evolve. Secondly, the web of affiliations connects not just agents, but agents and other entities such as knowledge, tasks and organizations. The agents described here in are more cognitively realistic than the typical a-life agents. They are also more socially realistic in terms of interaction than the typical e-commerce agents as the agents we use are boundedly rational rather than optimizers. Another distinction compared to most systems is that DyNet can take real networks as input.

In contrast to traditional SNA, DNA considers the role of the agent in terms of processes and not just position. That is, the agents can do things – communicate, store information, learn. Further, the networks are dynamic and changing even as the agents change. The links are probabilistic, the networks multi-colored and multi-plex to the extent that the set of networks combine in to one complex system where changes in one sub-network inform and constrain changes in the others, often leading to error cascades. Finally, DNA explores the sensitivity of the measures and the impacts to error.

The approach, theory, and results described here are illustrative. Clearly much work needs to be done before we have a complete understanding of network dynamics. Are there likely to be other change mechanisms than those currently in DyNet – to be sure. However, since all human action is cognitively mediated – it is unlikely that such mechanisms will not be derivable, at a basic level from what the physical and physiological constraints, what the agent knows, the basic learning and information processing mechanisms, and the way in which groups, organizations and institutions store such information. To create a truly dynamic network theory we need to create the equivalent of a quantum dynamics for the socio-cognitive world, where the fundamental entities, the people, unlike atoms, have the ability to learn.

References

Anderson, Brigham S., Carter Butts & Kathleen M. Carley, 1999, “The Interaction of Size and Density with Graph-Level Measures” to Social Networks. “ Social Networks. 21: 239-267.

Butts, C. 2000. “Network Inference, Error, and Informant (In)Accuracy: A Bayesian Approach,” to appear in Social Networks

Carley K. (1990). ``Group Stability: A Socio-Cognitive Approach,'' in Lawler E., Markovsky B., Ridgeway C. and Walker H. (Eds.), Advances in Group Processes}, JAI Press, Greenwich, CN.

Carley, Kathleen M. & Yuqing Ren, 2001, “Tradeoffs Between Performance and Adaptability for C3I Architectures.” In Proceedings of the 2001 Command and Control Research and Technology Symposium. Conference held in Annapolis, Maryland, June, 2001. Evidence Based Research, Vienna, VA.

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.” Greenwhich, CN: JAI Press, Inc. Stamford, CT, pp. 3-30.

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

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

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

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

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

Diedrich, Frederick J., Kathleen M. Carley, Jean MacMillan, Keith Baker, MAJ Jerry L. Schlabach, LTC J, Victor Fink, forthcoming, VISualization of Threats and Attacks (VISTA) in Urban Environments” Military Intelligence Professional Bulletin

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. Conference held in June. Monterray, CA. Evidence Based Research, Vienna, VA.

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. Conference held in June. Monterray, CA. Evidence Based Research, Vienna, VA.

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

Skvoretz, J. 1990. “Biased Net Theory: Approximations, Simulations, and Observations,” Social Networks, Vol. 12. pp. 217-238

Sparrow, M. (1991). The application of network analysis to criminal intelligence: An assessment of the prospects. Social Networks 13:251-274.