

Detecting Change in Longitudinal Social Networks

Ian McCulloh

Network Science Center, U.S. Military Academy, West Point, NY

ian.mcculloh@usma.edu

Kathleen M. Carley

Center for Computational Analysis of Social and Organizational Systems, School of Computer Science, Carnegie Mellon University

kathleen.carley@cs.cmu.edu

Abstract

Changes in observed social networks may signal an underlying change within an organization, and may even predict significant events or behaviors. The breakdown of a team's effectiveness, the emergence of informal leaders, or the preparation of an attack by a clandestine network may all be associated with changes in the patterns of interactions between group members. The ability to systematically, statistically, effectively and efficiently detect these changes has the potential to enable the anticipation, early warning, and faster response to both positive and negative organizational activities. By applying statistical process control techniques to social networks we can rapidly detect changes in these networks. Herein we describe this methodology and then illustrate it using four data sets, of which the first is the Newcomb fraternity data, the second set of data is collected on a group of mid-career U.S. Army officers in a week long training exercise, the third is the perceived connections among members of al Qaeda based on open source, and the fourth data set is simulated using multi-agent simulation. The results indicate that this approach is able to detect change even with the high levels of uncertainty inherent in these data.

Keywords

Statistical models for social networks, longitudinal social network analysis, Statistical Process Control, CUSUM, change detection

Acknowledgements

This research is part of the ARO Change Detection project with the USMA Network Science Center and the Dynamics Networks project in CASOS (Center for Computational Analysis of Social and Organizational Systems, <http://www.casos.cs.cmu.edu>) at Carnegie Mellon University.

This work was supported in part by:

- The Army Research Organization, MIPR No. 9FDATXR048.
- The Office of Naval Research (ONR), United States Navy Grant No. N00014-06-1-0104
- The Army Research Labs DAAD19-01-2-0009.
- The Army Research Institute ARI—W91WAW07C0063.
- Additional support on measures was provided by the National Science Foundation IGERT 9972762 in CASOS and the Department of Defense.

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 ARO, ONR, ARL, NSF, DOD or the U.S. government.

Introduction

Social network change detection (SNCD) represents an exciting new area of research. It combines the area of statistical process control and social network analysis. The combination of these two disciplines is likely to produce significant insight into organizational behavior and social dynamics. Immediate applications to counter terrorism and organizational behavior are possible due to the sheer volume of available electronic communications network data (McCulloh et al., 2008; Ring, Henderson & McCulloh, 2008).

Much research has been focused in the area of longitudinal social networks (Sampson, 1969; Newcomb, 1961; Romney et al., 1989; Banks & Carley, 1996; Sanil, Banks & Carley, 1995; Snijders, 1990, 2007; Frank, 1991; Huisman & Snijders, 2003; Johnson et al., 2003; McCulloh et al., 2007a, 2007b). Wasserman et al. (2007) state that, "The analysis of social networks over time has long been recognized as something of a Holy Grail for network researchers." Doreian & Stokman (1997) produced a seminal text on the evolution of social networks. In their book they identified as a minimum, 47 articles published in *Social Networks* that included some use of time, as of 1994. They also noted several articles that used over time data, but discarded the temporal component, presumably because the authors lacked the methods to properly analyze such data. An excellent example of this is the Newcomb (1961) fraternity data, which has been widely used throughout the social network literature. More recently, this data has been analyzed with its' temporal component (Doreian & Stokman, 1997; Krackhardt, 1998; Baller, et al. 2008).

Methods for the analysis of over-time network data have actually been present in the social sciences literature for quite some time (Katz & Proctor, 1959; Holland & Leinhardt, 1977; Wasserman, 1977; Wasserman & Iacobucci, 1988; Frank, 1991). Continuous time Markov chains for modeling longitudinal networks were proposed as early as 1977 by Holland & Leinhardt and by Wasserman. Their early work has been significantly improved upon (Wasserman, 1979; 1980; Leenders, 1995; Snijders & van Duijn, 1997; Snijders, 2001; Robins & Pattison, 2001) and Markovian methods of longitudinal analysis have even been automated in a popular social network analysis software package SIENA. A related body of research focuses on the evolution of social networks (Dorien, 1983; Carley (1990, 1991, 1995, 1999); Dorien & Stokman, 1997) to include three special issues in the *Journal of Mathematical Sociology (JMS 21, 1-2; JMS 25, 1; JMS 27, 1)*. Others have focused on statistical models of network change (Feld, 1997; Sanil, Banks, & Carley, 1995; Snijders, 1990, 1996; Van de Bunt et al., 1999; Snijders & Van Duijn, 1997). Robins & Pattison (2001, 2007) have used dependence graphs to account for dependence in over-time network evolution. We can clearly see that the development of longitudinal network analysis methods is a well established problem in the field of social networks.

We nominate four types of dynamic network behaviors for investigation in this paper. These behaviors are not comprehensive; however, it is necessary to define a set of behaviors to focus our investigation of network change. The four behaviors we focus our attention on include: network stability; endogenous change; exogenous change; and initiated change.

Stability occurs when the underlying relationship between agents in a network remains the same. It is possible that observed networks may contain error (Killworth & Bernard, 1976; Bernard & Killworth, 1977). If the network is stable, then changes in the network over time are due to observation error alone. An example of stability occurs in work environments where the underlying relationships remain unchanged, however, fluctuations exist as a result of stochastic noise, variations in daily work requirements, and sampling error.

Endogenous change occurs when the goals and motives of an individual, among other factors may drive the network to evolve. For example, a military platoon consisting of 20 to 30 soldiers can experience endogenous change as individuals interact, share beliefs and experiences. This is the focus of actor-

oriented models (Snijders, 2007) which attempt to estimate statistically significant behaviors, both structural and compositional, that drive network evolution. In a similar fashion, multi-agent simulation approaches attempt to investigate endogenous change by specifying agent-level behavior in order to infer network evolution.

Exogenous change occurs when a change is introduced separate from the agent interaction. With this type of change future events are independent from previous events. This implies that no inference can be drawn from the present model about the future network dynamics. An example of exogenous change might occur in the form of an enemy attack on a military platoon consisting of 20 to 30 soldiers. During the attack there is something fundamentally different about the relationships among the soldiers. There is nothing about the individual interactions that could predict this change caused by an exogenous source. In other situations, exogenous change can occur for many reasons. A shortage of economic resources could lead to job lay-offs that will significantly affect the social network, regardless of endogenous effects. These are of course drastic changes, presented here to illustrate abrupt forms of network change. It is also possible to have smaller change, such as when a new person joins a social group, a company finds new access to less expensive resources, or a group member finds a better way of accomplishing required tasks.

The final longitudinal network behavior we discuss is **initiated change**. We define this behavior as occurring when an exogenous change initiates a sequence of endogenous change. In our military example, it is possible that the heroic or cowardly actions of individuals in the platoon may affect the way other platoon members see them, thereby affecting the interaction among agents in the network and initiating endogenous network evolution.

It is important to delineate the difference between stability, endogenous, exogenous and initiated change if we are to understand network dynamics and any underlying processes governing network behavior. Again these changes are not comprehensive as one might imagine periodic change, event driven change, and other forms of change found in the dynamics literature. A first step toward the problem of longitudinal network analysis is to statistically determine that an organization has changed over time. For example, Johnson et al. (2003) studied people wintering over at the South Pole. There were three similar groups corresponding to three different years. A whole-network survey design was used to collect social network data once per month for eight months for each of the three groups. Johnson studied longitudinal change on the social networks of the three groups. Theoretically, these similar groups should exhibit similar evolutionary behavior. In one of the groups, there was an exogenous change that involved the “disappearance” of an expressive leader “due in part to harassment by a marginalized crew member.” This exogenous change significantly affected the evolutionary behavior of the network. This behavior was only apparent as a result of the similarity between the three groups and the large magnitude of the difference in network behavior, which enabled Johnson to determine the significant cause of this difference. In practice, this type of similarity among groups may be rare. SNCD offers a method to identify statistically significant abrupt change in network behavior in real-time, and to identify a likely change point of when the change occurred. This change point will allow a social scientist to identify potential causes of change, such as the disappearance of the crew member, and isolate that exogenous abrupt change from typical longitudinal behavior.

Our approach for detecting changes in longitudinal networks rapidly detects an abrupt change in some network measure over time. We are not predicting a future change, but rather rapidly identifying that a change has occurred; and then providing a statistically sound indication of when that change was likely to have occurred.

Rapid detection and identification of change is important for two key reasons. First, it allows an analyst monitoring a network in real time to respond quickly to organizational change, facilitating the change if it

is positive, and mitigating the effects of negative change on the organization. For example, ideas and policies are discussed and communicated within a network of people, long before organizational implementation. Sometimes, individual politics (network evolution) can prevent the implementation of good ideas (Rogers, 2003). Rapid detection of organizational change may cause a manager to investigate the presence of good initiatives and see them through to implementation. On the other hand, terrorist organizations will begin planning their attacks, long before they are actually carried out. Rapid change detection could alert military intelligence analysts to the shift in planning activities prior to the attack occurring.

The proposed approach may also be useful to social scientists investigating organizational change. This approach provides another tool for the exploration of longitudinal networks. Common problems with existing methods such as exponential random graphs and actor-oriented models include degeneracy and non-convergence (Handcock, 2003). SNCD can identify changes in longitudinal networks to help identify abrupt changes induced by some exogenous factor, such as the removal of the agent in the Johnson wintering over data (Johnson et al., 2003). With SNCD, the social scientist can identify shorter periods within the longitudinal network data where other methods may provide useful insight without convergence and degeneracy issues.

The third key reason that rapid change detection is important is that it limits the scope of explanation for network change. A sound statistical estimate of when a network change occurred can help a social scientist identify potential abrupt exogenous changes and thereby isolate periods of the network for more in-depth investigation. Determining the likely time of change in a network helps us understand where to look for fundamental conditions that cause groups to transform themselves. If we as social scientists could monitor networks in a daily or weekly basis, we could open a new line of research within longitudinal network analysis.

SNCD is essentially a statistical approach for detecting abrupt persistent changes in organizational behavior over time. Organizations are not static, and over time their structure, composition, and patterns of communication may change. These changes may occur quickly, such as when a corporation restructures, but they often happen gradually, as the organization responds to environmental pressures, or individual roles expand or contract. Often, these gradual changes reflect a fundamental qualitative shift in an organization, and may precede other indicators of change. It is important to note, however, that a certain degree of change is expected in the normal course of an unchanging organization, reflecting normal day-to-day variability. The challenge of Social Network Change Detection is whether metrics can be developed to detect signals of meaningful change in social networks in a background of normal variability.

This paper will introduce an application of statistical process control to detect change in longitudinal network data. A brief background is provided on statistical process control which is used extensively in manufacturing. Statistical process control is extended to social networks with important limitation and distribution assumptions being addressed. The newly proposed method is demonstrated on three longitudinal data sets. The performance of the method is then explored using multi-agent simulation.

Background

Longitudinal social network data is becoming increasingly more common. Longitudinal network data can be readily obtained in a semi-autonomous fashion from the internet, blogs, and email. Longitudinal network analysis is becoming increasingly relevant for the analysis of online citation networks, internet movie data, massive multi-player on-line games (MMPOG), patent data bases, phone-networks, email-based-networks, social-media networks and more.

Current methods of change detection in social networks, however, are limited. Hamming distance (Hamming, 1950) is often used in binary networks to measure the distance between two networks. Euclidean distance is similarly used for weighted networks (Wasserman & Faust, 1994). While these methods may be effective at quantifying a difference in static networks, they lack an underlying statistical distribution. This prevents an analyst from identifying a statistically significant change, as opposed to normal and spurious fluctuations in the network.

Jaccard indices are used by SIENA (Snijders et al., 2007) users to assess the amount of turnover from one observation of network panel data to the next. The amount of turnover may indicate a number of important features of the data, including whether an actor-oriented model is likely to have convergence issues. This index is not ideal for detecting network change for similar reasons as the Hamming distance.

The quadratic assignment procedure (QAP) and its multiple regression counterpart MRQAP (Krackhardt, 1987, 1992) has been used to detect structural similarity and compare networks in terms of their correlation. This is not the same as detecting a statistically significant change in the network over-time. The procedure could probably be adapted for such purpose, but this is not a trivial task and certainly beyond the scope of this paper.

Markovian approaches to longitudinal network analysis such as SIENA are good methods for modeling evolutionary change and determining structural factors that affect network change; however, these models may have convergence issues in the presence of sufficiently large abrupt endogenous or exogenous changes. These models also assume an underlying statistical process within the network that drives change, and models exogenous change with time dummies that requires some a priori knowledge of the change.

SNCD is a process of monitoring networks to determine when significant changes to their network structure occur so that analysts and researchers can more efficiently search for potential causes of change. We propose that techniques from social network analysis, combined with those from statistical process control can be used to detect when significant changes occur in longitudinal network data. In application, it requires the use of statistical process control charts to detect changes in observable network measures. By taking longitudinal measures of a network, a control chart can be used to signal when significant changes occur in the network. For those unfamiliar with statistical process control, it should be noted that the word “control” can be very misleading. In fact, nothing is controlled at all. Statistical process control is a collection of algorithms that monitor a stochastic process over time and rapidly detect statistically significant departures from typical behavior. Control charts refer to the individual algorithms used to monitor a process. The word “control” is derived from their application in quality control. Quality engineers attempt to control production lines by monitoring them and investigating any statistical anomalies. Through investigation, they attempt to mitigate negative process behavior and continue any newly discovered process improvements. In our application of SNCD, we use statistical process control to monitor longitudinal social networks and detect any statistically significant departures from typical behavior that may correspond to a change in the network. While the quality engineer uses this technique to “control” a manufacturing process, we envision that the social scientist will use it to gain insight in network dynamics.

There are many network measures that can be calculated from a given network. These include graph level measures, e.g., density, and node level measures, e.g., degree centrality. The SNCD technique is applicable to any measure of the network regardless of whether it is a graph level or a node level measure. In this paper for exposition purposes we focus on graph level measures rather than node level measures in order to investigate changes in the network as a whole as opposed to changes in the level of influence of a particular agent. For example, for each time period, we use the average of the betweenness (Freeman,

1977) over all nodes in the graph rather than the betweenness of a single node. The average betweenness may provide insight into group cohesion and the distribution of informal power throughout the organization. We also illustrate SNCD using density (Coleman & Moré, 1983), average closeness (Freeman, 1979), and average eigenvector centrality (Bonacich, 1972). Again, these measures provide slightly different insight into group cohesion. These four measures are chosen because they are commonly used in the literature and represent many potential measures available for change detection. Additional measures such as the maximum, minimum, and the standard deviation of the above node level measures are considered in a virtual experiment to explore limitations of the proposed method. A complete exploration of all social network measures and all possible types of changes to a network is certainly beyond the scope of this initial paper on the subject, however, we hope to have sufficiently illustrated the promise of this approach.

Another concern with these measures is their scale invariance. In order to compare measures across different time periods, they must be standardized. For a steady sized group this should not be an issue, but in the case of an expanding or contracting group, issues arise as to whether results can be used across the different scales of group size. In other words, the network measures may change in different ways with respect to the current group size and thus provide inconsistent information about the group even absent of any stochastic changes within the group. For more detailed information on the standardization of network measures, see Bonacich, Oliver & Snijders (1998). For this research, *ORA¹ developed by Kathleen Carley at the Center for Computational Analysis of Social and Organizational Systems at Carnegie Mellon University is used to compute the average network measures from all group information (Carley et al., 2009).

Statistical Process Control

SPC is a technique used by quality engineers to monitor industrial processes. They use control charts to detect changes in an industrial process by taking periodic samples from the process, calculating a statistic based on some process metric, and comparing the statistic against a decision interval. If the statistic exceeds the decision interval, the “control chart” is said to “signal” that a change may have occurred in the process. Once a potential change has been “signaled,” quality engineers investigate the process to determine if an actual change occurred, what the most likely time the change occurred was, and whether the process needs to be reset or improved to avoid financial loss for the company. Control charts are usually optimized for their processes to increase their sensitivity for detecting changes, while minimizing the number of “false positives”—signals when no change has actually occurred in the process.

Three control chart schemes are investigated in this paper; the cumulative sum (CUSUM) (Page, 1961); the Exponentially Weighted Moving Average (Roberts, 1959); and the Scan Statistic (Fisher & MacKenzie, 1922; Naus, 1965; Priebe et al., 2005). The CUSUM will be the primary method considered and recommended for longitudinal network analysis. This procedure provides an estimate of when the change actually occurred (change point detection) as opposed to simply signaling that a change occurred (change detection). The other two methods are applied to simulated networks in a virtual experiment to explore the performance of SNCD.

CUSUM

The CUSUM control chart (Page, 1961) was proposed as an improvement over the traditional Shewhart (1927) x-bar chart. The strength of the CUSUM was its use of sequential probability ratio testing which used information of previous observations to determine change in a stochastic process. Moustakides (2004) showed that the CUSUM procedure was a uniformly most powerful test for normally distributed processes with a specified size step change in the mean of the process. Unfortunately, in most applications

the investigator does not know a priori the size and type of the change. Furthermore, the underlying process may not be normally distributed. The quality engineering literature contains much exploration of the performance of the CUSUM under conditions of different magnitudes of change, types of change, and distributional assumptions.

The CUSUM control chart sequentially compares the statistic C_t against a decision interval h until $C_t > h$. Since one is not interested in concluding that the network process is unchanged, the cumulative statistic is

$$C_t^+ = \max\{0, Z_t - k + C_{t-1}^+\}$$

If this rule was not implemented the control chart would require more observations of the network to signal if $C_t < 0$ at the time of abrupt change. The statistic C_t^+ is compared to a constant, h^+ . If $C_t^+ > h^+$, then the control chart signals that an increase in the network measure might have occurred. In a similar fashion, $C_t^- = \max\{0, -Z_t - k + C_{t-1}^-\}$ and is compared to a constant, h^- . If $C_t^- > h^-$, then the control chart signals that a decrease in the network measure may have occurred.

To monitor for both directions of network change, two one-sided control charts are employed. One chart is used for monitoring increases in the monitored network property and the other is used for detecting decreases in the property. If the process remains in-control then C_t^\pm will fluctuate around zero. When $C_t^+ > h^+$ or $C_t^- > h^-$, the two one-sided CUSUM control chart scheme signals that the network may have changed.

Exponentially Weighted Moving Average Control Chart

The Exponentially Weighted Moving Average (EWMA) control chart was introduced by Roberts (1959) for monitoring changes in the mean of a process. The EWMA associated with subgroup t is

$$w_t = \lambda \bar{x}_t + (1 - \lambda)w_{t-1}, \text{ where } 0 < \lambda \leq 1 \text{ is the weight assigned to the current subgroup average and}$$

$$w_0 = \mu_0. \text{ Common values of } \lambda \text{ are } 0.1 \leq \lambda \leq 0.3. \text{ Having observed a total of } T \text{ subgroups, the statistic } w_T$$

is plotted against the decision interval $\mu_0 \pm L\sigma_{\bar{x}} \left(\frac{\lambda}{2 - \lambda} \left(1 - (1 - \lambda)^{2T} \right) \right)^{1/2}$, where L is a constant that scales the width of the decision interval.

Lucas & Saccucci (1987) (see also Saccucci & Lucas, 1990) investigated the impact of different combinations of L and λ on the average number of observations before the EWMA signals a change. The combinations that were investigated were chosen such that the false positive rate for each chart was the same. They found that EWMA charts with small values of λ perform well at detecting small changes in a process mean. Conversely, EWMA charts with large values of λ perform well at detecting large changes in a process mean. Hunter (1986) and Montgomery (1996) investigated the performance of the EWMA chart

and concluded that it is similar to the performance of the CUSUM chart. In addition, the EWMA is a time series approach for SPC. Therefore, the EWMA seems a good candidate for comparison to the CUSUM.

Scan Statistic

Scan statistics (Fisher & Mackenzie, 1922; Naus, 1965; Priebe, et. al., 2005), also known as *moving window analysis*, investigates a random field for the presence of a local signal. A small window of observations is used to calculate a local statistic. In this paper a window size of 7 observations proceeding the current time period is used, and the window mean is used for the local statistic. Increasing the window size reduces the likelihood of false alarm, but makes detection of a change less likely. Decreasing the window size makes the procedure more sensitive to change, but increases the probability of false signal. The decision to use a window size of 7 was chosen to be consistent with previous applications of the scan statistic for detecting longitudinal network changes (Priebe et al., 2005). If the statistic exceeds a decision interval, then inference can be made that a change in the network may have occurred.

Distributional Limitations

The performance and false alarm probability of the SPC procedures used in this approach assume that the stochastic process being monitored is independent and normally distributed. The assumptions are clearly violated in network applications. The degree to which these assumptions are violated and the impact on type I error varies based on the topology of the network. Networks that require a meaningful investment of resources to establish a link, limit the degree a node can obtain and the network tends to take on an Erdos-Renyi random topology (Erdos & Renyi, 1959; Alderson, 2009). In other networks, such as scale-free networks common for modeling the internet and certain biological networks, the distribution of many network measures is skewed and the false alarm rate may be adversely affected. Figure 1 shows the variance of data collected from a normal and right skewed distribution versus the number of observations sampled. The increased variance from the right skewed data will inflate the decision interval calculated on a few initial observations, making it more difficult to detect change, or more susceptible to false alarm.

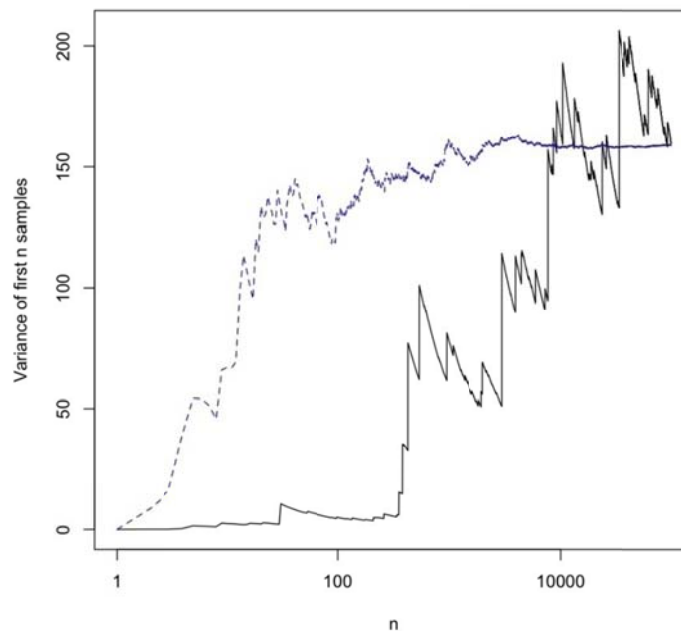


Figure 1. Bias Induced in Right Skewed Data

Some social scientists do not believe that groups can be adequately captured by quantitative analysis and statistical distributions (Brown & Morrow, 1994). We do not attempt to tackle this argument. Clearly, the work of this paper contributes to quantitative methods in social science. We also do not claim that a detected change is definitive proof that the organization has in fact changed. This approach will only detect a statistically significant change in the observed network measure of an organization. This could be a false alarm, an expected event affecting the organization, among other causes. Change detection simply alerts an analyst or social scientist that a change may have occurred. It is incumbent on the analyst or social scientist to investigate the group using many different methods in the social sciences to determine if change has in fact occurred, the nature of that change, and the cause of change. The approach laid out in this work will narrow the scope of this task by quickly identifying potential change and estimating when the change may have occurred.

Data

CUSUM is a method for assessing longitudinal change, and we use real-world data to demonstrate the practical application of the approach and simulated data to assess the accuracy of the approach. Altogether we use four data sets to demonstrate the efficacy of the social network change detection approach. We initially illustrate the CUSUM control chart on the Newcomb Fraternity data, a social network data set recorded of college transfer students; the Leavenworth data, a social network data set recorded of mid-career U.S. Army officers in a training exercise; and an al Qaeda data set. It is impossible to identify the “real” change in real-world data. For these data sets, we suggest compelling reasons for the change identified using SNCD; however, we acknowledge a different “story” might be constructed if different change points were identified. Thus, we also use simulated data generated by a multi-agent simulation so that we can decisively know the point of “real” change. Applying the CUSUM control chart to this data enables us to determine whether or not the proposed method can indeed identify the point of change. The performance comparison of the CUSUM to the EWMA, the Scan Statistic, and across various network level measures is explored using multi-agent simulation. The four data sets are explained in more detail.

Newcomb Fraternity Network

The first data set was collected by Theodore Newcomb (1961) at the University of Michigan. The participants included 17 incoming transfer students, with no prior acquaintance, who were housed together in fraternity housing. The participants were asked to rank their preference of individuals in the house from 1 to 16, where 1 is their first choice. Data was collected each week for 15 weeks, except for week number 9. David Krackhardt (1998) dichotomized the network data by assigning a link to preference ratings of 1-8 and having no link for ratings of 9 to 16. A visualization of the Newcomb Fraternity network for time period 8 is shown in Figure 2. The mean and standard deviation of the average betweenness, and average closeness was estimated from the first five networks to determine typical behavior. The CUSUM statistic was then calculated for all time periods. Note that the dichotomization scheme proposed by Krackhardt results in a constant density across all time periods, thus no change can occur in this measure.

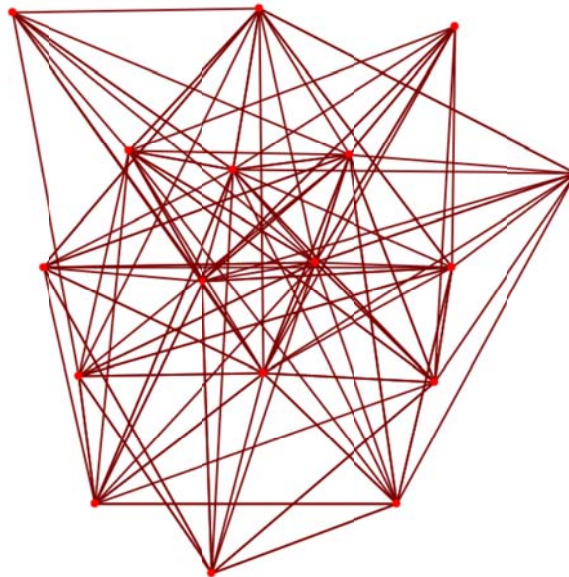


Figure 2. Dichotomized Newcomb Fraternity Network for Time Period 8.

Leavenworth Data

The second data set was collected from an Army war fighting simulation at Fort Leavenworth, Kansas, in April 2007, by Craig Schreiber. The participants were mid-career U.S. Army officers taking part in a brigade level staff training exercise. There were 68 participants in this data set, who served as staff members in the headquarters of the brigade conducting a simulated training exercise. Relational data was collected through self reported communications surveys over a period of four days, twice per day. Thus, there were 8 time periods. A directed relationship is recorded if an officer reports interacting with another one of the 68 officers during the preceding time period. Halfway through the second day (after time period 3), the brigade commander was displeased at the lack of coordination between the officers in the exercise. He brought all 68 participants together and chastised them for their performance and told them that they were expected to perform better. Therefore, SNCD might be able to indicate a significant change in the network corresponding to the brigade commander's interaction with the participants. This data set is unique in that it contains a known change point in time that can be used to validate the proposed method. Figure 3 shows the social network for time period 4 from the Leavenworth data set. The mean and standard deviation of the density, average betweenness, and average closeness was estimated from the first three networks to determine typical behavior. The CUSUM statistic was then calculated for all time periods. Three time periods were used because that represents about 30 percent of the time periods and is comparable to the number used with the Newcomb Fraternity data. Ideally, more networks will allow a more accurate estimate of typical behavior. The reader is reminded that these examples are used to illustrate the proposed methodology, while the performance of the method is evaluated using a simulated data set.

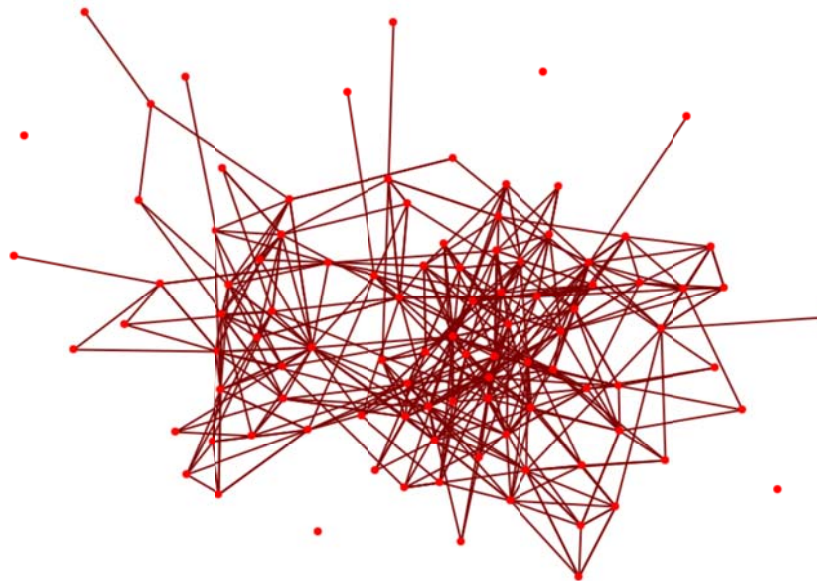


Figure 3. Leavenworth Network for Time Period 4

Al Qaeda Communications Network

The Center for Computational Analysis of Social and Organizational Systems (CASOS) at Carnegie Mellon University created snapshots of the annual communication between members of the al Qaeda organization from its founding in 1988 until 2004 from open source data (Carley, 2006). The data is limited in that we do not know the type, frequency, or substance of the communication and all links are non-directional, meaning we do not know who initiated communication with whom. Finally, the completeness of the data is uncertain since it only contains information available from open sources. The data is unique in that it provides a network picture of a robust network over standard time-periods of one year.

This data also provides a challenge for the proposed method due to the poor data quality. Bernard & Killworth (1979) state that “attempts at detecting change are useless unless data quality are high.” The fact that the proposed method succeeds at detecting change under these conditions speaks to its usefulness in practical applications.

Using the network snapshots for each year time-period, the average social network measures were calculated and plotted for betweenness, closeness, and density. Each of these measures increased from 1988 until 1994, and then leveled off. There are many possible reasons for this burn-in period, such as the quality of our intelligence gathering on al Qaeda and the rapid development and reorganization of a fast growing organization. In al Qaeda’s early years, access to the infant organization may have been limited, as well as the resources devoted to tracking a small, new, and relatively unaccomplished terrorist network. The organization itself may have also been changing drastically during its first years by actively recruiting new members, and shifting its structure to accommodate new resources and infrastructure.

A required condition for SNCD to be applied is a period of network stability. For this reason, the averages for each measure and standard deviation were calculated over the five years that follow the burn-in period that ended in 1994. The CUSUM control chart was then used to monitor the network from 1994 to 2004. Figure 4 is a snapshot of the al Qaeda social network.

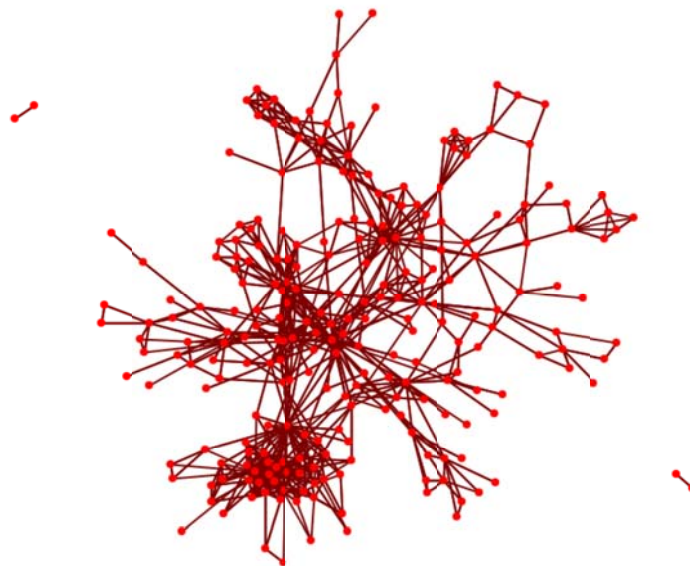


Figure 4. Monitored al Qaeda Communication Network for Year 2001

Simulated Data

Simulated data is used in order to inject an organizational change at a defined point in time. SNCD approaches can then be evaluated on their ability to identify that change. In real-world data, there are often many changes facing an organization and identifying one specific cause of change can be subjective or questionable. With simulated data, SNCD can be explored in a more controlled series of virtual experiments. For this initial investigation, we use a multi-agent simulation of a 100 node network, using the Construct² simulation model (Carley, 1990; Schreiber & Carley, 2004; Carley, Martin & Hirshman, 2009) set in the context of a U.S. infantry military organization (Headquarters, Department of the Army, 1992).

Construct is a dynamic-network multi-agent simulation grounded in constructivist theory (Carley, 1991; McCulloh et al., 2008). Agents are heterogeneous in their socio-demographic characteristics, information that they “know,” and their beliefs. Each time step agents may choose to interact with one or more others, communicate, and learn. The propensity of agents to interact is a function of knowledge, belief and task homophily; proximity of the agents; socio-demographic similarity, intent to learn new information, and intent to coordinate. Agent interaction leads to shared knowledge and thus greater knowledge-based homophily; however, heterophilous agents are less likely to interact. Construct has been validated in a number of settings and has been widely used to look at the co-evolution of social structure and culture, the diffusion of information and beliefs, and the impact of marketing campaigns and media on social behavior. Initial Construct populations, social and knowledge networks, can be hypothetical or real (Carley, Martin & Hirshman, 2009). Three key features that make Construct ideally suited to our needs are: 1) the social network evolves over time; 2) the user can specify “interventions” at specific times, thus guaranteeing a known state change in the system; and 3) the model can be instantiated with data on an actual group and so enables “what-if” reasoning about actual groups.

The basic military structure that was simulated was an infantry training model. This is the most basic U.S. military unit and is used for training soldiers and officers across the U.S. Army Training and Doctrine Command (Headquarters, Department of the Army, 1992). Within this model, soldiers are organized into four-man teams. Two teams and a squad leader form a 9-man squad. Three squads and a three-person headquarters form a 30-man platoon. Three platoons and a 10-person command post form a company. Each soldier is trained in various skills that are distributed throughout the organization. Each team, for example, will have an automatic gunner, a grenadier and two riflemen. One member on a team will also be trained as a medic, another in demolitions, and two will be able to search enemy prisoners of war. Each soldier possesses individual skill in stealth, situational awareness, physical fitness, intelligence, military rank, and motivation.

In the military context of this multi-agent simulation, the proximity was determined by the organizational proximity. Members of the same squad are closer to each other than other members in the platoon, who are closer than other members of the company. The socio-demographics of the agents do not change throughout the simulation and are coded as the agent’s military occupational specialty and military rank. The knowledge homophily was randomly seeded for each agent across 500 bits of knowledge data resulting in $3.27 * 10^{23}$ different agent knowledge combinations. This factor was allowed to change as agents share information when they interact, thus becoming more similar.

The simulation was verified by adjusting the relative weights applied to homophily, proximity, and socio-demographics. The model was validated, in 2008, by four military subject matter experts who confirmed that the simulated networks represent their experience of soldier relationships in military units.

The simulation was run with all agents present for the first 30 time periods. At time period 30, some type of change was imposed on the network, isolating some of the agents, thereby simulating radio failure or enemy attack. Figures 5 and 6 show example snapshots of the simulated network before and after the change.



Figure 5. Simulation before Change



Figure 6. Simulation after Change

The simulation was replicated 1,000 times to obtain estimates of the average time to detect change as well as the variance.

Method

Social network change detection algorithms are implemented in much the same way a control chart is implemented in a manufacturing process. Three different graph measures are used for change detection for the sake of illustrating the proposed method. SNCD can be applied to any node or graph measure over time. The graph measures for density, average closeness, and average betweenness centrality are calculated for several consecutive time-periods of the social network. The mean and variance for the measures of the network are calculated by taking a sample average and sample variance from networks that are assumed to be “typical.” At least two networks are required to estimate these values, however, more networks will allow a more accurate estimate of the mean and variance of the “typical” network measure. The subsequent, successive social network measures are then used to calculate the CUSUM’s C^+ and C^- statistics as well as the appropriate statistics for the EWMA and Scan Statistic. These are then compared to a decision interval to determine when or if the control chart signals a change in the mean of the monitored network measure. Upon receiving a signal, the change point is calculated by tracing the signaling C^+ or C^- statistic in the CUSUM procedure back to the last time period it was zero. In order to continue running the control chart after a signal, the mean and variance are recalculated after the network measures have stabilized following the change.

Recall that SNCD only indicates that a change may have occurred. The determination that the network has in fact changed and the subsequent determination that the network has stabilized following the change should be based on an investigation of other aspects of the network and the data surrounding the change point. Otherwise, the risk of misspecifying the change point can bias current and future findings of change.

This CUSUM methodology is demonstrated on three real-world data sets and explored in more detail through simulation. The real-world data sets are used to illustrate practical application of the approach. The decision threshold for the three real-world data sets was established at 3.0. If the network measure

were normally distributed, this would corresponded to an estimated risk of false alarm (type I error) of 0.01 (Galbreath, 2008). As noted earlier, as the distribution of the network measure is increasingly right skewed, bias is introduced that can increase the likelihood of false alarm. However, the network measures observed during the stabilized in-control period of the three data sets do not violate normality assumptions, as shown in the normal probability plots in Figure 7.

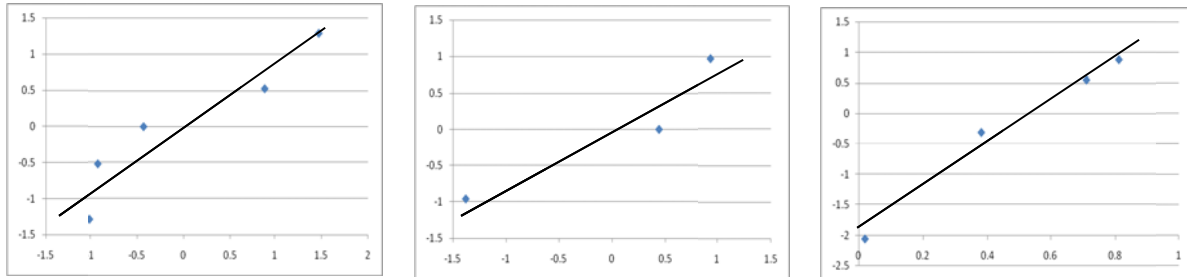


Figure 7. Normal Probability Plots of the In-Control Measures of Real-World Data

Virtual Experiment

A virtual experiment is conducted using the *Construct* Infantry Model to provide a realistic data set for evaluating SNCD methods. Three different size infantry units (squad, platoon, and company) are simulated for 500 time periods. In these units, four changes are introduced. This creates 9 independent data sets that can be used to evaluate SNCD performance. Three of the changes are not feasible for the squad size element. The four network changes correspond to common military communication problems that might affect an infantry unit.

The first type of network change is the isolation of the Headquarters section. For a squad, this is simply the squad leader. For a platoon, this consists of the platoon leader, platoon sergeant, and the radio telephone operator (RTO). For a company, this includes the 10-person command post, also known as the headquarters element. A military headquarters is most often isolated from the rest of the unit as a result of radio failure or a deliberate attack from enemy forces. This is perhaps one of the most significant changes that commonly happen in a military situation, as it requires a rapid and efficient transfer of command and control, as the formal hierarchy is significantly adjusted. In the simulation, this is modeled by isolating the headquarters section beginning at time period 20. These individuals remain isolated for the remainder of the simulation. Network measures are calculated on the organization for all time periods.

Another significant change in a military organization is the loss of a subordinate element. A subordinate element might be lost as a result of a task organization change, radio failure, or enemy attack. This change is not modeled for the infantry squad, since this would mean losing half of the organization. For the platoon, this change is modeled by isolating a squad at time period 20 for the remainder of the simulation. For the company, this is also modeled by isolating a squad at time period 20 for the remainder of the simulation. While it is conceivable to isolate any number of individuals in the simulation, these changes are used to demonstrate the performance of the SNCD methods. Perhaps SNCD methods that have similar performance could be evaluated under greater conditions of change in a future paper. For now, it is beyond the scope of this paper to exhaustively address all conceivable types of network change.

A similar change is the addition of a new subordinate element. This is usually a result of a task organization change. This is modeled by adding a squad in both the company and platoon level models. It

is not modeled for a squad, because squad organizations are not usually capable of managing an additional subordinate element. Again, this simple change is used to evaluate SNCD and not meant to be an exhaustive comparison of different types of organizational change.

The final type of change simulated, is sporadic communication. Sporadic communication can be either deliberate, or unplanned. An example of deliberate sporadic communication is a reconnaissance operation, where radio power must be conserved and noise discipline is important. An example of unplanned sporadic communication is radio failure. This is modeled in the simulation by introducing a squad from time period 30 to time period 40. Network measures will be recorded throughout the simulation. This change is only modeled for the platoon and company level simulations.

Table 1 illustrates the combinations of the virtual experiment. The outputs of the simulation are the graph level measures recorded for each simulated time step. Different SNCD methods are then used to identify possible changes in the network over time.

Table 1. Virtual Experiment

Variable	Number / Nature of Values	Values
Network Size	3	9, 30, 100
Type of Change in Network		
Isolation of leadership	2	Isolated headquarters after 30 time periods
Sporadic communication (reconnaissance)	2	Initially absent, present for 10 time periods, then absent for remainder of simulation (omitted for squad)
Loss of subordinate unit	2	Removal of the immediate subordinate unit after 30 time periods (omitted for squad)
Gain an attached unit	2	Addition of a squad after 30 time periods (omitted for squad)
Cells	18	3 Network sizes x 4 Changes x 2 Levels – Squad omissions
Replications	25	
Independent Runs	450	

The social network measures listed in Table 2 are measured for every simulated network.

Table 2. Social Network Measures

Average Betweenness	Standard Deviation of Closeness
Maximum Betweenness	Average Eigenvector Centrality
Standard Deviation of Betweenness	Maximum Eigenvector Centrality
Average Closeness	Minimum Eigenvector Centrality
Maximum Closeness	Standard Deviation of Eigenvector

Results

The approach proposed in this paper was found to be successful at detecting significant events in all data sets. Figure 8 displays a plot of the C statistics for Average Betweenness over time for the Newcomb Fraternity data. Recall that the CUSUM will detect either increases or decreases in a measure, but not both. Therefore, two control charts must be run for each social network measure monitored. In the figure, the two lines correspond to the chart for detecting increases in the measure and the chart for detecting decreases in the measure over time. The trends in the data for the betweenness measure are similar to the closeness measure. The density measure is not effective for change detection since the network is fixed-choice and the density remains 0.5 for every network.

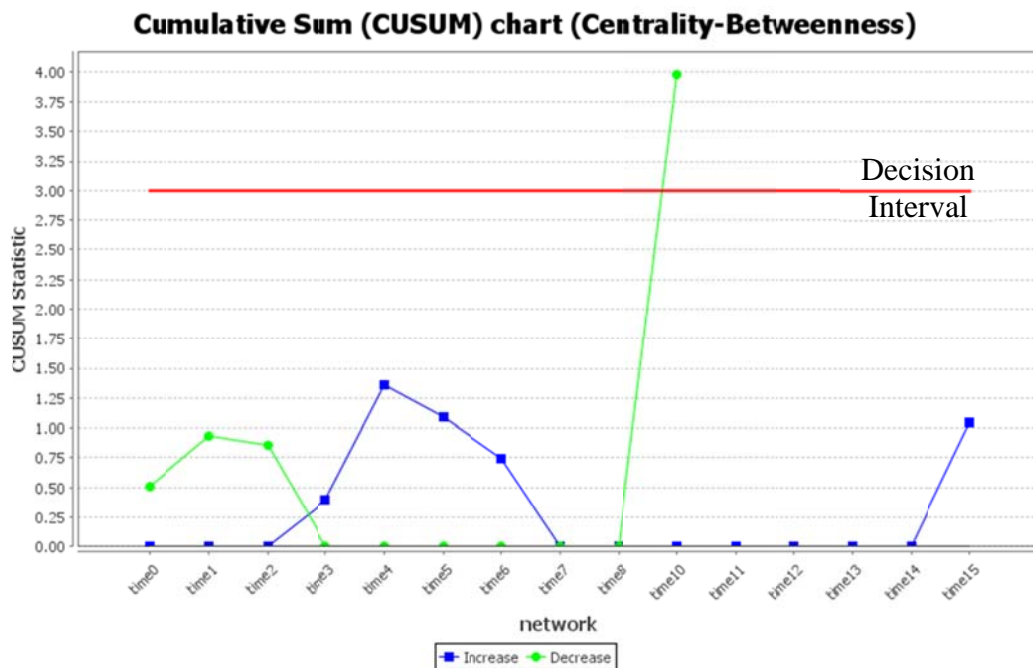


Figure 8. Plot of the CUSUM C Statistic Over Time for the Newcomb Fraternity Data

According to Figure 8, the control chart for average betweenness signals at time period 10 that a change may have occurred in the social network of the fraternity members. The most likely time that the change actually occurred is the last time period that the C statistic was equal to 0. This change point corresponds to time period 8 in the Newcomb Fraternity data, which was the week before a mid-semester break. It is not unreasonable that social relationships may have changed over a break, as participants possibly vacationed together. Unfortunately, the exact activities and dynamics of the group are not completely known. However, this data does provide evidence of the importance of the proposed method in analyzing network dynamics.

The Leavenworth data perhaps provides more compelling support for SNCD. Figure 9 illustrates the C statistics for average betweenness over time. The chart in Figure 9 signals at time period 5 that a change in the network may have occurred. The likely time the change actually took place is time period 3, which coincides with the brigade commander chastising the members of the group.

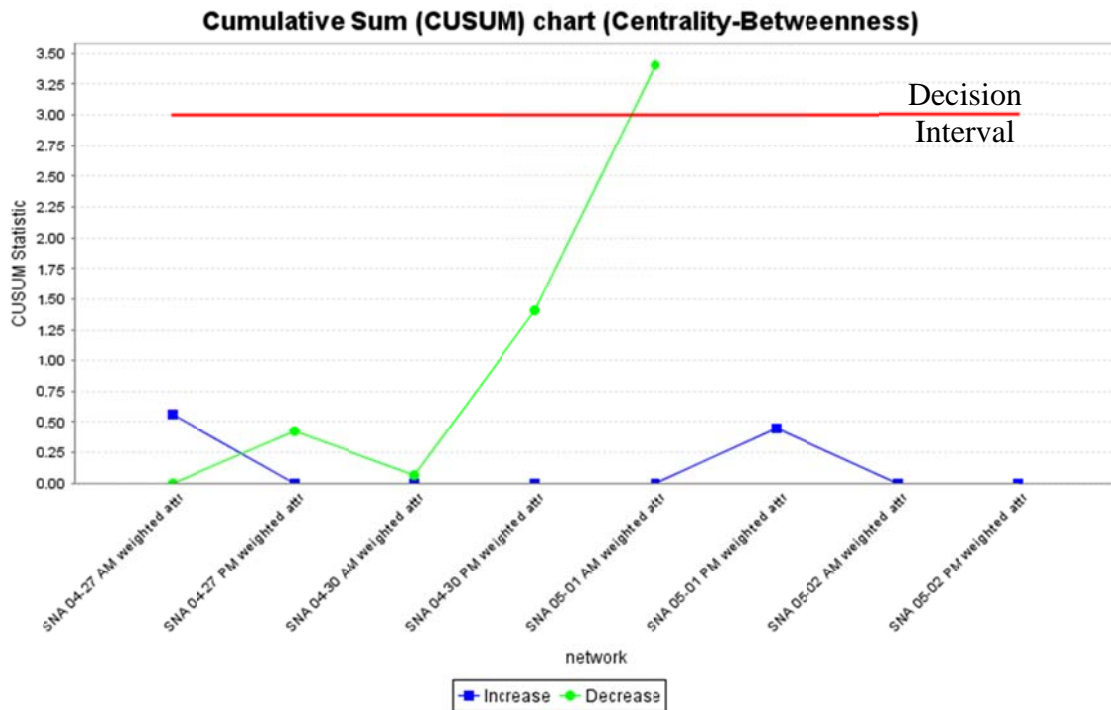


Figure 9. Plot of the CUSUM C Statistic Over Time for the Leavenworth Data

The al Qaeda data set offered data with more nodes that were aggregated over a much larger time period. At the same time, we were able to identify at least one major event in al Qaeda’s history. The question was asked, “Can we identify September 11 from the social network?” Perhaps more importantly, “Can we identify the point in time when the organization changed and began to plan the attacks?” Figure 10 shows the CUSUM statistic for the average betweenness of the al Qaeda network.

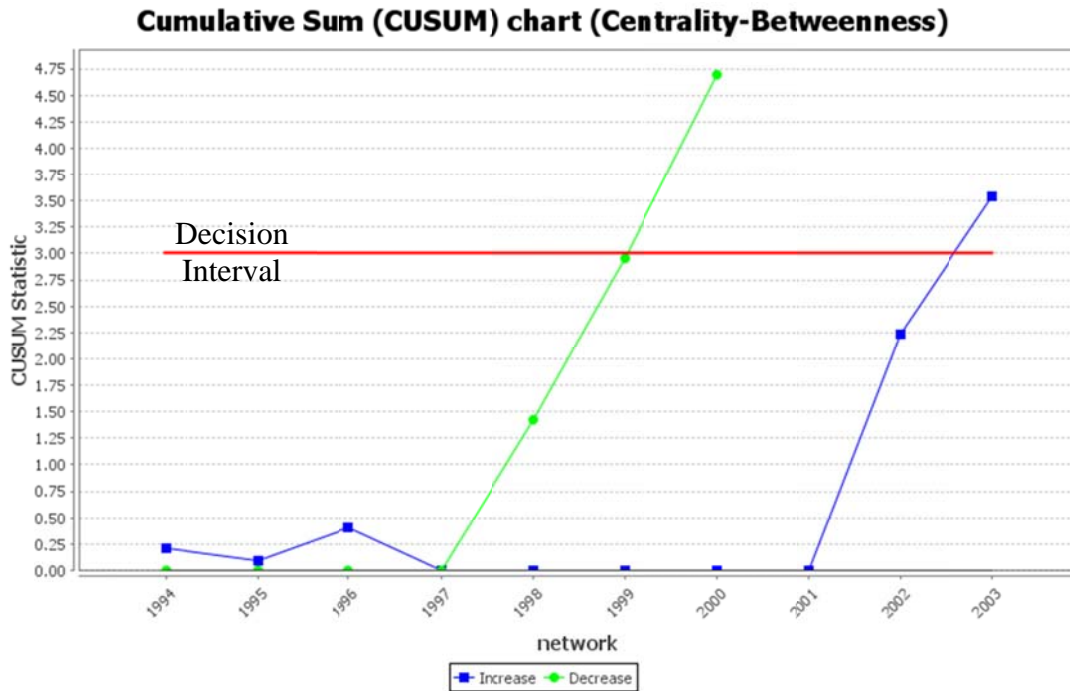


Figure 10. Plot of Betweenness CUSUM Statistic of al Qaeda

It can be seen in Figure 10 that the CUSUM statistic exceeds the decision interval and signals that there might be a significant change in the al Qaeda network, detected in the year 2000. Therefore, an analyst monitoring al Qaeda would be alerted to a critical, yet subtle change in the network prior to the September 11 terrorist attacks.

The CUSUM’s built in feature for determining the most likely time that the change occurred estimates the change point as 1997. For the density and closeness measures, this point in time is also 1997. To understand the cause of the change in the al Qaeda network, an analyst should look at the events occurring in al Qaeda’s internal organization and external operating environment in 1997.

Several very interesting events related to al Qaeda and Islamic extremism occurred in 1997. Six Islamic militants massacred 58 foreign tourists and at least four Egyptians in Luxor, Egypt (Jehl, 1997). United States and coalition forces deployed to Egypt in 1997 for a bi-annual training exercise were repeatedly attacked by Islamic militants. The coalition suffered numerous casualties and shortened their deployment. In early 1998, Zawahiri and Bin Laden were publicly reunited, although based on press release timing, they must have been working throughout 1997 planning future terrorist operations. In February 1998, an Arab newspaper introduced the “International Islamic Front for Combating Crusaders and Jews.” This organization established in 1997, was founded by Bin Laden, Zawahiri, leaders of the Egyptian Islamic Group, the Jamiat-ul-Ulema-e-Pakistan, and the Jihad Movement in Bangladesh,

among others. The Front condemned the sins of American foreign policy and called on every Muslim to comply with God's order to kill the Americans and plunder their money (Marquand, 2001). Six months later the US embassies in Tanzania and Kenya were bombed by al Qaeda. Thus, 1997 was possibly the most critical year in uniting Islamic militants and organizing al Qaeda for offensive terrorist attacks against the United States. It is interesting that the proposed SNCD method identifies and accurately determines when change occurred.

Virtual Experiment Results

Using the social simulation program, *Construct* (Carley, 1990; Carley, 1995; Schrieber & Carley, 2004), the performance of SNCD was explored through simulation. A variety of changes are introduced to the network at a known point. The Cumulative Sum (CUSUM), Exponentially Weighted Moving Average (EWMA), and Scan Statistic, statistical process control charts are applied to several social network graph level measures taken on the network at each time step. The number of time steps between the actual change and the time that an SNCD method "signals" a change will be recorded as the Detection Length. The Average Detection Length (ADL) over multiple independently seeded runs is then a measure of the SNCD method's performance. The ADL will be compared for different changes and different SNCD parameters.

Isolation of Headquarters

Investigating the isolation of the headquarters element in three different organizations will provide insight into how the network size affects the performance of change detection measures. In each organization (30-man platoon, 100-man company, and 9-man squad); 10 percent of the network was removed. In a sense, the magnitude of change is the same; however, the network size is different.

The isolation of the platoon headquarters is modeled by removing the three headquarters members at time period 30 for the duration of the simulation. Social network measures are recorded for all time periods. Table 3 displays the ADL performance of the SNCD methods. It can be seen that the average of the betweenness is a better measure to use for SNCD than either the maximum or the standard deviation of betweenness. This is generally true for all magnitudes of change and sizes of organization investigated. For the closeness measure, both the maximum closeness and average closeness generally outperform the standard deviation of closeness. However, for an EWMA with $r = 0.3$, the maximum closeness measure has relatively poor performance. This might suggest that the average closeness measure is a more robust measure of change detection. In a single variant, non-network application of the EWMA, the parameter, r , makes the control chart more or less sensitive to a particular magnitude of change (Lucas & Saccucci, 1990; McCulloh, 2004). It is reasonable to consider that for the isolation of a platoon headquarters, the maximum closeness EWMA with $r \leq 0.2$ is sensitive to detecting the change, yet the maximum closeness EWMA with $r \geq 0.3$ is less sensitive. This will be explored with other magnitudes and types of changes throughout the paper. For eigenvector centrality, the maximum eigenvector centrality and the standard deviation of eigenvector centrality appear to be more sensitive measures of change detection than the average or minimum of the eigenvector centrality. It also appears that the eigenvector centrality measures dominate all other measures for performance in this case.

Table 3. ADL Performance of SNCD on Isolation of Platoon Headquarters

	CUSUM $k = 0.5$	EWMA $r = 0.1$	EWMA $r = 0.2$	EWMA $r = 0.3$	Scan Statistic
Average Betweenness	9.32	8.24	10.16	11.52	6.76
Maximum Betweenness	14.36	14.72	15.72	17.08	13.24
Std. Dev. Betweenness	16.44	16.24	16.92	18.52	15.24
Average Closeness	10.68	9.08	13.60	17.52	10.48
Maximum Closeness	8.76	6.00	10.60	37.96	8.64
Std. Deviation Closeness	34.48	34.72	34.52	35.68	27.08
Average Eigenvector	31.28	31.28	31.28	31.28	24.00
Minimum Eigenvector	14.36	14.36	14.28	15.56	14.88
Maximum Eigenvector	5.24	5.40	5.80	7.52	4.00
Std. Dev. Eigenvector	5.92	4.88	6.40	6.96	3.64

Statistical process control is a powerful statistical method for detecting the change. Figure 11 shows four measures plotted for the same simulated longitudinal networks. The top two plots are the network measure of betweenness over time. The bottom two plots are the CUSUM statistic C calculated on the same betweenness measure over time. The two plots on the left show the measures plotted when there is no change present in the network over time. These plots show stochastic fluctuations induced by the simulation. The two plots on the right show the measures plotted when a change is imposed at time period 20. The change is identified much more clearly using the CUSUM, especially when the reader directs their attention to the scale of the y-axis in the four plots.

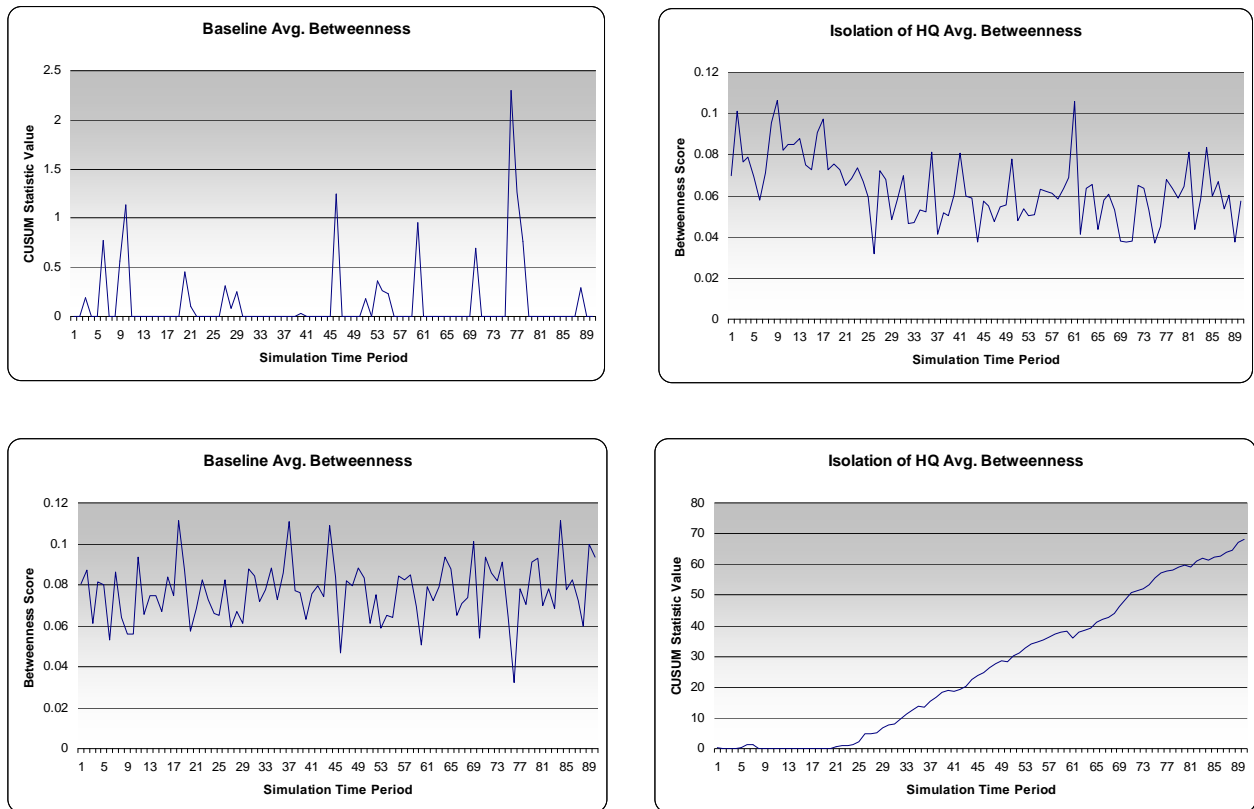


Figure 11. Plots of the Average Betweenness Centrality (top) Compared to Plots of the CUSUM Statistic, C (bottom) for Situations with No Change (left) and with Change (right)

The visual identification other types of change imposed on the network, and other SNCD schemes yield similar success. The CUSUM is simply used to illustrate the power of the general change detection approach. Other magnitudes and types of change will be compared by simply reporting the ADL from when a change occurs until the SNCD scheme signals.

The isolation of the company headquarters was modeled by removing the 10 soldier headquarters section at time 30 for the remainder of the simulation. This is very similar to the platoon example, in that 10 percent of the organization is removed. Social network measures are again recorded for all time periods. Table 4 displays the ADL performance of each of the SNCD methods applied to the 100 node network. Again, it can be seen that the average of the betweenness is a more effective measure of change detection than the maximum or the standard deviation of betweenness. The performance of the closeness measures behave as they did in the case of platoon headquarters isolation. In this case, the maximum eigenvector centrality does not appear to be as effective of a measure for detecting change as does other measures. However, the standard deviation of eigenvector centrality still dominates all other measures for change detection performance.

Table 4. ADL Performance of SNCD on Isolation of Company Headquarters

	CUSUM $k = 0.5$	EWMA $r = 0.1$	EWMA $r = 0.2$	EWMA $r = 0.3$	Scan Statistic
Average Betweenness	11.16	11.08	10.20	13.48	6.96
Maximum Betweenness	17.32	17.76	18.20	20.12	13.72
Std. Dev. Betweenness	18.08	19.40	20.88	22.52	17.36
Average Closeness	11.16	9.44	12.52	15.64	9.40
Maximum Closeness	10.44	9.72	12.64	51.76	9.60
Std. Deviation Closeness	41.88	39.48	42.20	43.44	40.76
Average Eigenvector	35.84	36.72	34.84	34.84	29.24
Minimum Eigenvector	16.00	17.96	17.88	16.76	13.60
Maximum Eigenvector	26.40	30.76	29.64	29.24	25.44
Std. Dev. Eigenvector	10.40	10.72	9.36	9.48	6.44

The isolation of squad leadership was modeled by removing the squad leader at time 20 for the remainder of the simulation. This is also similar in that 11 percent of the organization is isolated. Table 5 shows the SNCD performance at the squad level, 9 node network. It is not clear that certain measures perform better than others for change detection in the 9 node network. It appears that the measures of average betweenness, average closeness, and the standard deviation of eigenvector centrality become better measures of network change as the size of the network increases. However, they do not necessarily perform worse on a small network. While an extensive study of the sensitivity of each measure to the network size is beyond the scope of this paper, it holds the promise of fruitful future research.

Table 5. ADL Performance of SNCD on Isolation of Squad Leader

	CUSUM <i>k</i> = 0.5	EWMA <i>r</i> = 0.1	EWMA <i>r</i> = 0.2	EWMA <i>r</i> = 0.3	Scan Statistic
Average Betweenness	16.12	15.76	16.32	17.92	12.32
Maximum Betweenness	16.64	17.40	19.52	18.56	11.56
Std. Dev. Betweenness	17.68	17.76	18.20	18.72	12.08
Average Closeness	15.16	15.84	16.48	15.60	11.72
Maximum Closeness	18.72	19.60	18.68	23.80	14.32
Std. Deviation Closeness	16.20	16.08	15.52	16.24	12.88
Average Eigenvector	24.12	24.12	24.12	24.12	15.12
Minimum Eigenvector	17.84	18.48	17.04	18.08	12.36
Maximum Eigenvector	19.36	21.56	20.56	20.56	13.84
Std. Dev. Eigenvector	17.08	18.72	18.36	17.44	12.36

Loss of Subordinate Element

The loss of a subordinate element provides insight into how the magnitude of change affects change detection performance. For the 30 man platoon and the 100 man company, a nine man squad is isolated. This represents 30 percent of the platoon and 9 percent of the company. This change is obviously not feasible for the nine man squad, since it would involve removal of the entire organization.

The infantry platoon had one squad removed from the simulation at time period 20, for the remainder of the simulation. Social network measures were recorded for each time period. The ADL for each measure is reported in Table 6. Again, it can be seen that the average of the betweenness outperforms other betweenness measures. The closeness measures perform as in previously investigated cases. The minimum eigenvector centrality outperforms the maximum eigenvector centrality for most of the SNCD schemes for this particular type and magnitude of change. The standard deviation of eigenvector centrality still outperforms other eigenvector centrality measures, however, it is no longer dominates all other measures.

Table 6. ADL Performance for Loss of Subordinate Element in a Platoon

	CUSUM $k = 0.5$	EWMA $r = 0.1$	EWMA $r = 0.2$	EWMA $r = 0.3$	Scan Statistic
Average Betweenness	6.96	6.00	8.68	12.16	8.12
Maximum Betweenness	9.52	7.44	11.12	13.24	7.80
Std. Dev. Betweenness	9.16	7.40	9.48	12.72	6.84
Average Closeness	9.64	8.36	12.72	19.28	11.40
Maximum Closeness	9.32	9.16	12.36	31.56	9.52
Std. Deviation Closeness	18.96	16.44	19.40	26.24	17.04
Average Eigenvector	29.36	29.36	29.36	29.36	20.60
Minimum Eigenvector	10.08	9.64	12.24	12.60	10.28
Maximum Eigenvector	11.72	12.04	11.88	20.60	10.84
Std. Dev. Eigenvector	8.48	6.28	9.80	10.44	6.88

The infantry company also had one squad removed at time 20 for the remainder of the simulation. The results for the company network are shown in Table 7. It generally takes longer to detect the changes in the company network. This was also observed in the isolation of the headquarters. This implies that the size of the network could impact the speed of change detection. The average betweenness, average closeness, and standard deviation of eigenvector centrality appear to outperform other measures for change detection performance. The maximum closeness measure dominates other measures in all cases except for the EWMA with $r = 0.3$.

Table 7. ADL Performance for Loss of Subordinate Element in a Company

	CUSUM $k = 0.5$	EWMA $r = 0.1$	EWMA $r = 0.2$	EWMA $r = 0.3$	Scan Statistic
Average Betweenness	13.64	11.72	13.80	20.60	12.68
Maximum Betweenness	23.80	19.64	23.80	30.72	25.44
Std. Dev. Betweenness	24.84	18.12	24.96	25.52	22.04
Average Closeness	9.72	7.4	13.44	14.96	9.80
Maximum Closeness	6.92	4.92	7.48	53.16	6.32
Std. Deviation Closeness	45.44	47.92	47.96	50.88	43.68
Average Eigenvector	34.72	36.60	34.72	34.72	30.64
Minimum Eigenvector	18.68	19.96	19.64	23.88	18.32
Maximum Eigenvector	18.28	25.80	25.00	27.20	25.88
Std. Dev. Eigenvector	9.52	9.92	11.88	15.32	8.72

Addition of New Subordinate Element

Another type of change is the addition of a new subordinate element. A squad is added to both the 30-man platoon and the 100-man company.

The infantry platoon had one squad that was not present initially, and added at time period 20. Social network measures were calculated for each time period. SNCD methods were applied to the data. Results are shown in Table 8. Although the speed of change detection is much faster for this type of change, the same performance trends are seen as before. For betweenness measures, the average outperforms the maximum or the standard deviation. The average closeness and maximum closeness measure perform well, however, the maximum closeness does not perform well with an EWMA $r = 0.3$ scheme. The standard deviation of eigenvector centrality almost completely dominates other measures.

Table 8. ADL Performance for Addition of Subordinate Element in a Platoon

	CUSUM $k = 0.5$	EWMA $r = 0.1$	EWMA $r = 0.2$	EWMA $r = 0.3$	Scan Statistic
Average Betweenness	1.60	1.52	1.68	1.72	1.00
Maximum Betweenness	2.32	2.16	2.20	2.00	1.00
Std. Dev. Betweenness	2.36	2.36	2.40	2.24	1.00
Average Closeness	1.48	1.52	1.56	1.52	1.00
Maximum Closeness	1.24	1.28	1.20	5.00	1.00
Std. Deviation Closeness	3.44	4.60	4.20	3.48	2.64
Average Eigenvector	31.76	31.76	31.76	31.76	25.56
Minimum Eigenvector	6.24	5.6	6.16	6.80	4.20
Maximum Eigenvector	4.52	4.88	4.80	4.80	3.56
Std. Dev. Eigenvector	1.16	1.60	1.24	1.24	1.00

The company model had a squad added at time period 20 for the remainder of the simulation. Again the platoon level performance is better than the company level performance, shown in Table 9. The average betweenness, average closeness, and maximum closeness all perform well at detecting the change. Surprisingly, the standard deviation of eigenvector centrality is not an effective measure for this type and magnitude of change.

Table 9. ADL Performance for Addition of Subordinate Element in a Company

	CUSUM $k = 0.5$	EWMA $r = 0.1$	EWMA $r = 0.2$	EWMA $r = 0.3$	Scan Statistic
Average Betweenness	9.64	9.52	9.84	10.28	5.04
Maximum Betweenness	14.52	16.96	15.80	17.44	12.16
Std. Dev. Betweenness	12.88	13.16	13.32	14.56	8.92
Average Closeness	5.32	5.8	5.36	5.24	1.44
Maximum Closeness	4.24	5.12	4.48	6.04	1.04
Std. Deviation Closeness	10.40	18.52	12.96	12.32	10.00
Average Eigenvector	35.56	37.04	38.64	37.60	30.24
Minimum Eigenvector	38.16	39.32	38.04	40.84	36.40
Maximum Eigenvector	30.20	33.48	34.44	29.52	30.92
Std. Dev. Eigenvector	33.88	33.72	37.80	44.48	33.96

Sporadic Communication

Sporadic communication was modeled with a squad communicating from time period 30 to time period 40 only. It can be seen in Table 10 that the performance of different measures is much more similar than in previous types of change. It is also interesting that all of the ADL values are greater than 10, which means that the change was detected after the organization returned to its original state. This might be a result of the SNCD statistic being moved closer to the decision interval from time period 30 to time period 40. When the organization returned to its original state, the statistic is much closer to the decision interval than it was before the change occurred. Therefore, the statistic is much more likely to signal a false positive after the sporadic change than it is to detect an actual change. This increased sensitivity can therefore provide an alert that a sporadic change may have occurred.

Table 10. ADL Performance for Sporadic Communication

	CUSUM <i>k</i> = 0.5	EWMA <i>r</i> = 0.1	EWMA <i>r</i> = 0.2	EWMA <i>r</i> = 0.3	Scan Statistic
Average Betweenness	15.08	14.20	16.12	17.56	17.76
Maximum Betweenness	15.24	16.52	16.88	18.24	17.84
Std Dev. Betweenness	14.28	14.80	16.04	17.40	17.48
Average Closeness	13.72	13.68	16.84	16.80	17.52
Maximum Closeness	12.44	12.16	15.32	18.32	17.20
Std Deviation Closeness	23.16	19.96	21.76	21.36	17.24
Average Eigenvector	24.32	24.32	24.32	24.32	18.84
Minimum Eigenvector	12.76	14.32	11.92	12.80	14.56
Maximum Eigenvector	12.96	12.68	14.36	14.36	18.84
Std. Dev Eigenvector	12.88	14.20	16.80	16.48	21.28

All methods of SNCD were ineffective for detecting sporadic changes in the company network. The sporadic change did not persist long enough to signal a possible change in most of the runs. The squad level network was not investigated for this type of change, due to a lack of context.

Conclusion

Statistical process control is a critical quality-engineering tool that provides rapid detection of change in stochastic processes (Montgomery, 1991; Ryan, 2000). The three real-world examples and the virtual experiments presented in this paper demonstrate that SNCD could enable analysts and researchers to detect important changes in longitudinal network data. Furthermore, the most likely time that the change occurred can also be determined. This allows one to allocate minimal resources to tracking the general patterns of a network and then shift to full resources when changes are determined.³ SNCD is therefore, an important analysis method for studying network dynamics.

It is critical to be able to detect change in networks over time and to determine when observed fluctuations are not simply stochastic noise. This paper describes a method for change detection based off of statistical process control, and then demonstrates its ability to detect changes in networks. Within this method, three specific control chart schemes for detecting change were considered: CUSUM, Exponentially Weighted Moving Average, and a Scan Statistic. No doubt other change detection methods will emerge and control chart schemes will emerge.

We found the CUSUM technique to be robust and to be of value in applied settings. The strengths of the proposed method are its statistical approach, its utility with a wide range of social network metrics, its ability to identify change points in organizational behavior, and its flexibility for various magnitudes of change. The proposed method requires the assumption of a period of stability that is necessary to estimate the mean and standard deviation of social network measures for “typical” network observations. In addition, the proposed method requires a reasonable number of time periods in which to detect change; i.e., greater than four.

The empirical results described in this paper, such as the detection of change in the al Qaeda network should be viewed with caution. We present them here purely to illustrate the methodology. Limitations on the data make it difficult to determine the validity of the results; thus, we should simply view these results as showing the promise of this methodology. The Leavenworth data spans only four days and used self-reported survey data, therefore it is not likely that it captured all communication and interaction among officers. The fact that even in this data set we were able to systematically detect a key change suggests the value of the proposed approach. The al Qaeda data, was based on open source information. As such it is an incomplete representation of interaction in that terror network. We cannot be sure that we have the entire communication network, or even a true picture of the observed communication network. However, the fact that our technique detects a change corresponding with the 9/11 attacks is intriguing. This work suggests that our approach may provide some ability to detect change even when there is incomplete information.

That being said, it is important that future work examine the errors associated with this technique, both the false positives and false negatives. Future work should also consider the sensitivity of this approach to missing information, and to the reason why the information is missing. For example, data sets collected post-hoc that focus on activity around an event, such as the al Qaeda data are prone to errors of missing nodes and as a result links prior to the event. In addition, open-source data tends to over-focus on nodes whose centrality is assumed; often resulting in “popular” actors being possibly over-connected and less popular actors being under-connected. Whereas, data sets collected based on opportunity, such as the Leavenworth data, are prone to missing links among the nodes.

In order to rectify the above shortcomings, future research should focus on improved methods for node and link inference or near-complete datasets with high resolution. Higher resolution involves taking many snapshots of the network. This may mean, simply an increase in frequency, e.g. changes by month, or it

may mean a longer time horizon, e.g., more years. The right choice will depend on the problem where we want to detect network change. More data points will provide more opportunities to detect changes while they are still small, instead of allowing them to incubate and grow as was the case for the al Qaeda data. As a minimum two observed networks are required to estimate the “typical” behavior of a social group being monitored for change. In practice, five or more networks are preferred to reduce the variance in estimating the statistical process control parameters. Larger datasets will also provide near continuous network measures permitting the use of control charts for continuous data. Near complete data means that the data should cover the communication network, with little or no missing information for a large contiguous period. Here one might consider simply tracking a group in general, as opposed to focusing on tracking relative to a specific event. Data such as that on the U.S. Congress or Supreme Court that is regularly output might provide a good source of data.

Another limitation of this approach is that the over-time dependence assumptions are ignored. This is common in statistical process control. English et al. (2001) points out that “the independence assumption is dramatically violated in processes subjected to process control.” Many manufacturing processes include feedback control systems which create autocorrelation among factors affecting the process. This is similar to problems of dyadic dependence and ergodicity issues with networks. In practice however, statistical process control still provides a great deal of insight, identifying when a process changes. This is no different in a network application. Networks may even have less dependence issues than manufacturing processes. Most manufacturing processes are engineered with feedback and control in an attempt to optimize the process. This is not necessarily true with social networks. Robins and Pattison (2007) lay out several statistical tests involving dependence graphs that can be used to determine if dependence is a statistically significant problem in a network. Just like the issues of normality, the dyadic dependence in the network can be verified similar to residual analysis in regression. If dependence is an issue in the network, SNCD can still be used to determine that a change occurred, however, there may be bias and an increase in the probability of a false positive. Future research should investigate both the impact of dependence on ADL performance as well as methods to better handle the problem statistically.

Social networks may also exhibit periodicity over time. Intuitively, people’s communication patterns may change in cycles over time. People tend to communicate with different people during the week, while at work, than on the weekends. People may communicate more frequently at certain times of the day. Even seasonal trends may affect observed social networks. The application of wavelet theory and Fourier analysis in particular may provide insight into the periodic behavior of network dynamics. Methods should be developed to test and filter periodicity from network measures over time. This will allow SNCD to be more accurate in determining the time a change actually occurred and may reduce the ADL for certain changes.

Future research should also look at the sensitivity of the optimality constant, k and control limit values of the CUSUM control chart for network measure change detection. As stated earlier, these values are generally arbitrarily chosen and then optimized for the process. By using further Monte Carlo simulations, a researcher should determine which parameter value would be best in detecting certain types of changes such as sudden large changes or slow creeping shifts. Usage of control charts on comparing models and observations should also be studied to see what specific conclusions can be obtained.

Multi-agent simulations provide valuable insight into the performance of control charts for social network change detection applications. Simulations allow an investigator to introduce various changes into a simulated organization and evaluate the time to detect for different algorithms. Simulations provide an efficient means of evaluating change detection on social networks. More importantly, however, is the ability to create more controlled experiments, by fixing certain variables, exploring others, and using

many replications to estimate error. Simulation studies will continue to be extremely useful in exploring extensions of this methodology.

Social network change detection is important for identifying significant shifts in organizational behavior. This provides insight into policy decisions that drive the underlying change. It also shows the promise of enabling predictive analysis for social networks and providing early warning of potential problems. In the same way that manufacturing firms save millions of dollars each year by quickly responding to changes in their manufacturing process, social network change detection can allow senior leaders and military analysts to quickly respond to changes in the organizational behavior of the socially connected groups they observe. The combination of statistical process control and social network analysis is likely to produce significant insight into organizational behavior and social dynamics. As a scientific community we can hope to see more research in this area as network statistics continue to improve.

References

- Alderson, D. (2009). "Catching the 'Network Science' Bug: Insight and Opportunities for the Operations Researchers," *Operations Research* 56, 5: 1047–1065.
- Baller, D., J. Lospinoso & A.N. Johnson (2008). "An Empirical Method for the Evaluation of Dynamic Network Simulation Methods." In *Proceedings of The 2008 World Congress in Computer Science Computer Engineering and Applied Computing*, Las Vegas, NV.
- Banks, D.L., & K.M. Carley (1996). "Models for Network Evolution." *Journal of Mathematical Sociology* 21: 173-196.
- Bernard, H.R. & P.D. Killworth (1977). "Informant Accuracy in Social Network Data II." *Human Communications Research* 4: 3-18.
- Bonacich, P. (1972). "Factoring and Weighting Approaches to Clique Identification." *Journal of Mathematical Sociology* 2: 113–120.
- Bonacich, P., A. Oliver & T.A.B. Snijders (1998). "Controlling for Size in Centrality Scores." *Social Networks* 20, 2: 135-141.
- Brown, R.A. & D.D. Morrow (1994). *Critical Theory and Methodology*. Thousand Oaks, CA: Sage.
- Carley, K.M. (1990). "Group Stability: A Socio-Cognitive Approach." *Advances in Group Processes* 7: 1-44.
- Carley, K.M. (1991). "A Theory of Group Stability." *American Sociology Review* 56, 3: 331–354.
- Carley, K.M. (1995). "Communication Technologies and Their Effect on Cultural Homogeneity, Consensus, and the Diffusion of New Ideas." *Sociological Perspectives* 38, 4: 547-571.
- Carley, K.M. (1999). "On the Evolution of Social and Organizational Networks." *Research in the Sociology of Organizations* 16: 3-30.
- Carley, K.M. (2006). "A Dynamic Network Approach to the Assessment of Terrorist Groups and the Impact of Alternative Courses of Action." In *Visualising Network Information Meeting Proceedings RTO-MP-IST-063*. Neuilly-sur-Seine, France: RTO. Available: http://www.vistg.net/documents/ISTO63_PreProceedings.pdf [January 7, 2011].

- Carley, K.M., J. Reminga, J. Storricks, & M. De Reno (2009). *ORA User's Guide 2009*. Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report CMU-ISR-09-115. Available: <http://www.casos.cs.cmu.edu/publications/papers/CMU-ISR-09-115.pdf> [January 7, 2011].
- Carley, K.M., M.K. Martin & B. Hirshman (2009). "The Etiology of Social Change," *Topics in Cognitive Science* 1, 4.
- Coleman, T. F. & J.J. Moré (1983). "Estimation of Sparse Jacobian Matrices and Graph Coloring Problems." *SIAM Journal on Numerical Analysis* 20, 1: 187–209.
- Doreian, P. (1983). "On the Evolution of Group and Network Structures II: Structures within Structure." *Social Networks* 8: 33-64.
- Doreian, P. & F.N. Stokman (Eds.) (1997). *Evolution of Social Networks*. Amsterdam: Gordon and Breach.
- English, J.R., T. Martin, E. Yaz & E. Elsayed (2001). "Change Point Detection and Control Using Statistical Process Control and Automatic Process Control." Presentation at the IIE Annual Conference, 2001, Dallas, TX.
- Erdős, P. & A. Rényi (1959). "On Random Graphs I." *Publicationes Mathematicae* 6: 290–297.
- Feld, S. (1997). "Structural Embeddedness and Stability of Interpersonal Relations." *Social Networks* 19: 91-95.
- Fisher, R.A., H. Thornton & W. Mackenzie (1922). "The Accuracy of the Plating Method of Estimating the Density of Bacterial Populations, with Particular Reference to the Use of Thornton's Agar Medium with Soil Samples." *Annals of Applied Biology* 9: 325–359.
- Frank, O. (1991). "Statistical Analysis of Change in Networks." *Statistica Neerlandica* 45: 283–293.
- Freeman, L. (1977). "A Set of Measures of Centrality Based on Betweenness." *Sociometry* 40: 35-41.
- Freeman, L. (1979). "Centrality in Social Networks I: Conceptual Clarification." *Social Networks* 1: 215-239.
- Hamming, R.W. (1950). "Error Detecting and Error Correcting Codes." *Bell System Technical Journal* 26, 2:147-160.
- Handcock, M. S. (2003). "Assessing Degeneracy in Statistical Models of Social Networks." Working Paper No. 39. Center for Statistics and the Social Sciences, University of Washington. Available: <http://www.csss.washington.edu/Papers/wp39.pdf> [January 7, 2011].
- Headquarters, Department of the Army (1992). *Field Manual 7-8, Infantry Rifle Platoon and Squad*. U.S. Army Infantry School, Ft. Benning, GA.
- Holland, P. & S. Leinhardt (1977). "A Dynamic Model for Social Networks." *Journal of Mathematical Sociology* 5, 5-20.

- Huisman, M., & T.A.B. Snijders (2003). "Statistical Analysis of Longitudinal Network Data with Changing Composition." *Sociological Methods and Research* 32: 253-287.
- Hunter, J.S. (1986). "The Exponentially Weighted Moving Average." *Journal of Quality and Technology* 18: 203-210.
- Jehl, D. (1997). "Islamic Militants Attack Tourists in Egypt." *The New York Times*, November 23, 1997. p. WK2.
- Johnson, J.C., J.S. Boster & L.A. Palinkas (2003). "Social Roles and the Evolution of Networks in Extreme and Isolated Environments." *Journal of Mathematical Sociology* 27: 89-121.
- Katz, L. & C.H. Proctor (1959). "The Configuration of Interpersonal Relations in a Group as a Time-Dependent Stochastic Process." *Psychometrika* 24: 317-327.
- Killworth, P.D. & H.R. Bernard (1976). "Informant Accuracy in Social Network Data." *Human Organization* 35:269-286.
- Krackhardt, D. (1987). "QAP Partialling as a Test of Spuriousness." *Social Networks* 9: 171-186.
- Krackhardt, D. (1992). "A Caveat on the Use of the Quadratic Assignment Procedure." *Journal of Quantitative Anthropology* 3: 279-296.
- Krackhardt, D. (1998). "Simmelian Tie: Super Strong and Sticky." In R. Kramer & M. Neale (Eds.), *Power and Influence in Organizations*. Thousand Oaks, CA: Sage, 21-38.
- Leenders, R. (1995). "Models for Network Dynamics: A Markovian Framework." *Journal of Mathematical Sociology* 20: 1-21.
- Lucas, J.M. & M.S. Saccucci (1990). "Exponentially Weighted Moving Average Control Schemes: Properties and Enhancements." *Technometrics* 32: 1-12.
- Marquand, R. (2001). "The Tenets of Terror." *Christian Science Monitor*, October 18, 2001.
- McCulloh, I., G. Garcia, K. Tardieu, J. MacGibbon, H. Dye, K. Moores, J.M. Graham & D.B. Horn (2007). *IkeNet: Social Network Analysis of Email Traffic in the Eisenhower Leadership Development Program*. (Technical Report, No. 1218). Arlington, VA: U.S. Army Research Institute for the Behavioral and Social Sciences.
- McCulloh, I., J. Lospinoso & K.M. Carley (2007). "Social Network Probability Mechanics." In *Proceedings of the World Scientific Engineering Academy and Society 12th International Conference on Applied Mathematics*, Cairo, Egypt, December 29-31, 2007.
- McCulloh, I., B. Ring, T. Frantz, & K.M. Carley (2008). "Unobtrusive Social Network Data from Email." In *Proceedings, 26th Army Science Conference*. Orlando, FL, December 1-4, 2008.
- McCulloh, I. (2004). *Generalized Cumulative Sum Control Charts*. Master's Thesis, The Florida State University.
- Montgomery, D.C. (1991). *Introduction to Statistical Quality Control*, 2nd edition. New York: John Wiley and Sons.

- Moustakides, G.V. (2004). "Optimality of the CUSUM Procedure in Continuous Time." *Annals of Statistics* 32, 1: 302-315.
- Naus, J. (1965). "Clustering of Random Points in Two Dimensions." *Biometrika* 52: 263-267.
- Newcomb, T.N. (1961). *The Acquaintance Process*. New York: Holt, Rinehart and Winston.
- Page, E.S. (1961). "Cumulative Sum Control Charts." *Technometrics* 3: 1-9.
- Priebe, C.E., J.M. Conroy, D.J. Marchette & P. Youngser (2005). "Scan Statistics on Enron Graphs." *Computational and Mathematical Organization Theory* 11: 229-247.
- Ring, B., S. Henderson & I. McCulloh (2008). "Gathering and Studying Email Traffic to Understand Social Networks." In H.R. Arabnia & R.R. Hashemi (Eds.), *Proceedings of the 2008 International Conference on Information and Knowledge Engineering, IKE 2008*, July 14-17, 2008. Las Vegas, NV: CSREA Press, 338-343.
- Roberts, S.V. (1959). "Control Chart Tests Based on Geometric Moving Averages." *Technometrics* 1: 239-250.
- Robins, G. & P. Pattison (2001). "Random Graph Models for Temporal Processes in Social Networks." *Journal of Mathematical Sociology* 25: 5-41.
- Robins, G. & P. Pattison (2007). "Interdependencies and Social Processes: Dependence Graphs and Generalized Dependence Structures." In: P. Carrington, J. Scott & S. Wasserman (Eds.), *Models and Methods in Social Network Analysis*. New York: Cambridge University Press, 192-214.
- Rogers, E.M. (2003). *Diffusion of Innovations*, 5th edition. New York, NY: Free Press.
- Romney, A.K. (1989). "Quantitative Models, Science and Cumulative Knowledge." *Journal of Quantitative Anthropology* 1: 153-223.
- Ryan, T. P. (2000). *Statistical Methods for Quality Improvement*, 2nd edition. Wiley.
- Saccucci, M.S. & J.M. Lucas (1990). "Average Run Lengths for Exponentially Weighted Moving Average Control Schemes Using the Markov Chain Approach." *Journal of Quality Technology* 22: 154-159.
- Sampson, S.F. (1969). *Crisis in a Cloister*. Ph.D. Thesis, Ithaca, NY: Cornell University.
- Sanil, A., D. Banks & K.M. Carley (1995). "Models for Evolving Fixed Node Networks: Model Fitting and Model Testing." *Social Networks* 17, 1: 65-81.
- Schreiber, C. & K.M. Carley (2004). *Construct; A Multi-agent Network Model for the Co-Evolution of Agents and Socio-Cultural Environments*. Carnegie Mellon University, School of Computer Science, Institute for Software Research International, Technical Report, CMU-ISRI-04-109. Available: <http://reports-archive.adm.cs.cmu.edu/anon/isri2004/CMU-ISRI-04-109.pdf> [January 7, 2011].
- Shewhart, W.A. (1927). "Quality Control." *Bell Systems Technical Journal* 6, 4 (October 1927): 722-735.

- Snijders, T. A. B., & M.A.J. Van Duijn (1997). "Simulation for Statistical Inference in Dynamic Network Models." In R. Conte, R. Hegselmann & P. Tera (Eds.), *Simulating Social Phenomena*. Berlin: Springer, 493-512.
- Snijders, T.A.B. (1990). "Testing for Change in a Digraph at Two Time Points." *Social Networks* 12: 539-573.
- Snijders, T.A.B. (1996). "Stochastic Actor-Oriented Models for Network Change." *Journal of Mathematical Sociology* 21: 149-172.
- Snijders, T.A.B. (2001). "The Statistical Evaluation of Social Network Dynamics." In: Sobel, M.E. & M.P. Becker (Eds.), *Sociological Methodology*. Boston: Basil Blackwell, 361-395.
- Snijders, T.A.B. (2007). "Models for Longitudinal Network Data." In: P. Carrington, J. Scott & S. Wasserman (Eds.), *Models and Methods in Social Network Analysis*. New York: Cambridge University Press, 148-161.
- Snijders, T.A.B., C.E.G. Steglich, M. Schweinberger & M. Huisman (2007). *Manual for SIENA version 3.1*. University of Groningen: ICS/Department of Sociology; University of Oxford: Department of Statistics. Available: http://stat.gamma.rug.nl/sie_man31.pdf [January 7, 2011].
- Van de Bunt, G.G., M.A.J. Van Duijn & T.A.B. Snijders (1999). "Friendship Networks through Time: An Actor-Oriented Statistical Network Model." *Computational and Mathematical Organization Theory* 5: 167-192.
- Wasserman, S. (1977). *Stochastic Models for Directed Graphs*. Ph.D. dissertation, Harvard University, Department of Statistics, Cambridge, MA.
- Wasserman, S. (1979). "A Stochastic Model for Directed Graphs with Transition Rates Determined by Reciprocity." In K.F. Schuessler (Ed.), *Sociological Methodology*. San Francisco: Jossey-Bass, 392-412.
- Wasserman, S. (1980). "Analyzing Social Networks as Stochastic Processes." *Journal of American Statistical Association* 75: 280-294.
- Wasserman, S. (2007). "Introduction." In P.J. Carrington, J. Scott, & S. Wasserman (Eds.), *Models and Methods in Social Network Analysis*. New York: Cambridge University Press.
- Wasserman, S. & D. Iacobucci (1988). "Sequential Social Network Data." *Psychometrika* 53, 2: 261-282.
- Wasserman, S., & K. Faust (1994). *Social Network Analysis: Methods and Applications*. New York: Cambridge University Press.

¹ *ORA can be downloaded from <http://www.casos.cs.cmu.edu/projects/ora/> [January 7, 2011].

² Construct is available at <http://www.casos.cs.cmu.edu/projects/construct> [January 7, 2011].

³ Three social network change detection algorithms (Shewhart X-Bar, Cumulative Sum, and Exponentially Weighted Moving Average) are available in the "Statistical Network Monitoring Report" in the software tool, Organizational Risk Analyzer (ORA) available through the Center for Computational Analysis of Social and Organizational Systems (CASOS), <http://www.casos.cs.cmu.edu> [January 7, 2011].