DYNAMIC NETWORK CHANGE DETECTION

MAJ Ian McCulloh
United States Military Academy
West Point, New York 10996

Professor Kathleen M. Carley
Center for Computational Analysis of Social and Organizational Systems
Institute for Software Research, School of Computer Science
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

ABSTRACT

Network data provides valuable insight into understanding complex organizations by modeling relational dependence between network agents. Detecting subtle changes in organizational behavior can alert analysts before the change significantly impacts the larger group. Statistical process control is applied to dynamic network measures of longitudinal data to quickly detect organizational change. The performance of 10 network measures and three algorithms are evaluated on simulated data. One of the algorithms and one of the network measures are used to demonstrate change detection on the Al-Qaeda terrorist network. There is no statistically significant difference in the performance of investigated algorithms, however, the cumulative sum control chart has a built-in estimate of the actual time a change may have occurred.

1. INTRODUCTION

According to the National Research Council's Committee on Network Science for Future Army Applications (2005), Network Science is defined as, "*the study of network representations of physical, biological, and social phenomena leading to predictive models of these phenomena. Initiation of a field of network science would be appropriate to provide a body of rigorous results that would improve the predictability of the engineering design of complex networks and also speed up basic research in a variety of applications areas.*" An important step forward in improving the predictability of networks is Dynamic Network Change Detection. Underlying changes in physical, biological, and social phenomena may be observed in the networks of these "application areas". "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 of change, provide early warning of change, and enable faster response to change." (McCulloh and Carley, n.d.)

Statistical process control is applied to dynamic network measures in a novel approach to the statistical analysis of networks. The area of statistical process control and quality engineering is as well established in academia as dynamic network analysis itself. The combination of these two disciplines is likely to produce significant insight into organizational behavior and social dynamics. Immediate applications to counter terrorism are obvious. There is also evidence that dynamic network change detection can enhance the command and control of friendly military units.

Dynamic network change detection is applied to longitudinal observed network data to rapidly detect small persistent changes in the underlying structure being modeled. We assume that these structures are not fixed and that their relationships, attributes and composition may change over time. These changes may be gradual or rapid. The changes may represent normal network evolution or they may represent a fundamental shift or shock in the application area being modeled. While some degree of variability is expected in sequential observations of an unchanged network, the challenge of dynamic network change detection is to detect significant change in a background of noise.

This paper will provide a brief background of longitudinal network analysis in network science. An overview of three dynamic network change detection methods is presented; the cumulative sum (CUSUM), the exponentially weighted moving average (EWMA), and a scan statistic (SS). Statistical process control is extended to dynamic networks and demonstrated on three longitudinal data sets. Findings suggest that the cumulative sum control chart is a good method for detecting changes in network behavior. Dynamic network change detection represents an exciting new area of research which will significantly push the frontier of network science.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
01 DEC 2008	**N/A**	**-**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Dynamic Network Change Detection	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **United States Military Academy West Point, New York 10996**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	I0. SPONSOR/MONITOR'S ACRONYM(S)
	I1. SPONSOR/MONITOR'S REPORT NUMBER(S)

I2. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited

I3. SUPPLEMENTARY NOTES
See also ADM002187. Proceedings of the Army Science Conference (26th) Held in Orlando, Florida on 1-4 December 2008, The original document contains color images.

I4. ABSTRACT

I5. SUBJECT TERMS

I6. SECURITY CLASSIFICATION OF:			I7. LIMITATION OF ABSTRACT	I8. NUMBER OF PAGES	I9a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**UU**	**6**	

2. BACKGROUND

Longitudinal network data may indicate different relationships at different points in time. While there may exist an underlying relationship that does not change, such as friendship, co-workers, mentorship, familial ties, the actual observance of a link may or may not be present at different points in time. Intuitively, every individual has a set of friends or acquaintances. The individuals do not communicate with all friends and acquaintances at all times. Therefore, the absence of a link in a single network does not mean that there is not a relationship between the nodes. It is also possible for a link to be present when there is no relationship. Examples include junk email, inadvertently hitting "reply to all" in email, or typing in a web address incorrectly. This is an important concept in understanding the longitudinal behavior of networks.

Several methods have been proposed to model longitudinal networks. Exponential random graph models (ERGM) have been used (Snjiders, 2005; Schweinberger, 2005). A link probability model (LPM) has been shown to resemble empirical data better than the ERGM in certain studies (McCulloh and Carley, 2008a; Baller, Lospinoso and Johnson, 2008). A multi-agent simulation based on constructural theory can add relational dependence into the LPM, creating more realistic longitudinal networks (McCulloh, Lospinoso, and Carley, 2008). Several methods for longitudinal analysis of networks have even been automated in software (Carley et al 2008; Snijders, 2008). These methods present a useful framework, where methods of statistical process control can be investigated for their efficacy in detecting changes in longitudinal networks.

Statistical process control (SPC) is a technique used by quality engineers to monitor industrial processes. They use SPC to detect changes in important quality characteristics by taking periodic samples from the process, calculating a statistic, and comparing the result against a decision interval. If the statistic exceeds the decision interval, a change in the process may have occurred and it is investigated. There are two risks for financial loss. One risk is in not detecting the process change quick enough, creating financial loss for the company by making substandard or wasteful product. The other risk is in stopping the process to search for a change that doesn't exist; a false alarm. SPC methods are usually optimized for their specific processes to increase their sensitivity for detecting changes, while minimizing the risk of false alarms.

Three common SPC methods that we consider here are the CUSUM (Page, 1961), EWMA (Roberts, 1959), and the SS (Fisher and Mackenzie, 1922). These methods are used in quality engineering to detect small changes in a process (Montgomery, 1991; Ryan, 2000). Larger changes may be more obvious, therefore, we restrict our attention to identifying small subtle changes in network behavior.

In this paper we posit that a network is ultimately the result of a stochastic process governing link formation within a network. SPC can be applied to measurements of a network in the same way it can be applied to measurements of quality characteristics of items produced on a manufacturing floor. In this manner, various measurements of a network will be used to calculate an SPC statistic and compared to a decision interval for significant change. A representative sample of network measures, SPC methods, magnitudes of change, and size of network is presented. A complete listing of these results would exceed the available space for this paper. For a more comprehensive report of change detection results, the reader is referred to McCulloh and Carley (2008a, 2008b).

3. METHOD

Dynamic network change detection is applied to network metrics in the same way statistical process control charts are applied to quality metrics in a manufacturing process. Ten different network level measures are investigated for change detection performance, along with three different SPC methods using simulated data. When the network measures appear to have stabilized over time, the "in-control" mean and variance for the measures of the network are calculated by taking a sample average and sample variance of the stabilized measures. The subsequent, successive dynamic network measures are then used to calculate the statistics for the CUSUM, the EWMA, and the SS. These are then compared to decision intervals to determine when or if the method signals a change in the mean of the monitored network measure. Each of the SPC methods is calibrated to have the same sensitivity to false alarms using the simulated data. The sensitivity to false alarm is 0.01, which corresponds to a false alarm every 100 observations on average. In order to continue running the control chart if a false alarm should occur, the in-control mean and variance of the monitored network measure are recalculated after the network measure has stabilized following the change.

Network data is simulated using *Construct*, a multi-agent simulation based on constructural theory (Carley, 1990; Carley 1995; Schrieber and Carley, 2004). The simulation models an Army Infantry company and has been validated and docked against C3TRACE and IMPRINT simulation models (McCulloh and Carley, 2008a). Nodes are isolated in the network at a known point in time. The CUSUM, EWMA, and SS SPC methods are applied to the 10 network level measures

taken on the network at each time step. The number of time steps between the actual change and the time that an SPC 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 SPC method's performance. The ADL will be compared for different magnitudes of change and for different network measures across three SPC methods.

Once the ADL performances of the three SPC methods have been explored, they will all be applied to a real world data set. The real world data set will consist of network data on the Al-Qaeda terrorist network as reported by McCulloh, Carley, and Webb (2007).

4. RESULTS

Certain network measures yield better change detection performance than others. Ten different network measures were investigated for use in network change detection. The ADL demonstrates the network measures' relative performance compared to other network measures. A lower ADL suggests that SPC run on that measure is more likely to detect a network change faster than using other network measures. The ADL results of the ten different network measures are displayed in Table 1. For this particular change, the 10 person headquarters from the 100 person Infantry company was isolated in the simulation at time period 30. The ADL reports the average number of networks that were observed after time period 30, until the SPC method signaled that a change may have occurred in the network.

Table 1. ADL Performance of Network Measures.

	CUSUM	EWMA	SS
Avg Betweenness	***11.16***	***11.08***	***9.96***
Max Betweenness	17.32	17.76	13.72
St Dev Betweenness	18.08	19.40	17.36
Avg Closeness	***11.16***	***9.44***	***9.40***
Max Closeness	***10.44***	***9.72***	***9.60***
St Dev Closeness	41.88	39.48	40.76
Avg Eigenvector Cent	35.84	36.72	29.24
Min Eigenvector Cent	16.00	17.96	13.60
MaxEigenvector Cent	26.40	30.76	25.44
St Dev Eigenvector Cent	***10.40***	***10.72***	***6.44***

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 average and maximum of the closeness measures outperform the standard deviation of closeness. Finally, the standard deviation of eigenvector centrality also demonstrates excellent change detection properties.

Changes were investigated for networks of size 9, 30, and 100, corresponding to an Infantry squad, platoon, and company. Various sizes of change were also explored to include isolating 1, 3, 10, 19, 30, and 40 nodes. Changes where nodes were added to the network were also explored. The ADL trends of the measures were relatively consistent across all of the measures. In many cases the maximum closeness measure did not perform as well as the other successful measures. For a more comprehensive presentation of the ADL performance of measures, see McCulloh and Carley (2008a).

A reader should be cautioned against comparing the SPC methods across the rows of Table 1. There is no statistically significant difference in the values across rows. For example, the ADL for maximum betweenness under the CUSUM is reported as 17.32, which is better than under the EWMA at 17.76. However, for the same network change, the ADL for average betweenness is 11.16 under the CUSUM, which is worse than the 11.08 under the EWMA. There is a significant difference between the ADLs of the network measure average betweenness which is approximately 11 and maximum betweenness which is approximately 17 observations.

The performance of the algorithms were investigated by comparing the ADL of each SPC method over increasing magnitudes of change. The magnitude of change presented here involve isolating 1, 3, 10, 19, and 30 nodes corresponding to the company commander, command group, company headquarters, a squad with the company headquarters, and a platoon respectively. A random removal of nodes shows similar behavior, however, we have chosen to present the scenario based results here. Table 2 shows the ADL of the CUSUM, EWMA, and SS applied to the average betweenness for different magnitudes of change. The column on the left indicates the number of nodes removed from the network at time 30. The ADL is reported as the average number of networks observed after time period 30 until the SPC method signals that a change may have occurred.

Table 2. ADL performance for magnitude of change.

Change	CUSUM	EWMA	Scan
0	106.13	103.85	107.27
1	88.32	86.52	96.2
3	22.49	21.34	25.51
10	11.32	8.04	12.28
19	2.44	1.72	1.60
30	1.12	1.28	1.00

It can be seen in Table 2 that the ADLs for all SPC methods are much smaller as more nodes are isolated from the network. This demonstrates the method's success at detecting change. Figure 1 displays a plot of the data in Table 2. There is no statistically significant difference in ADL performance for the three SPC methods for small changes of three nodes or less. For the

10 node change, the EWMA outperforms the CUSUM, which outperforms the SS. For the 19 node change the SS outperforms the EWMA which outperforms the CUSUM. For the 30 node change the EWMA has the worst overall performance. These findings are consistent with the performance of these methods in single variable, independent and identically distributed applications (McCulloh, 2004).

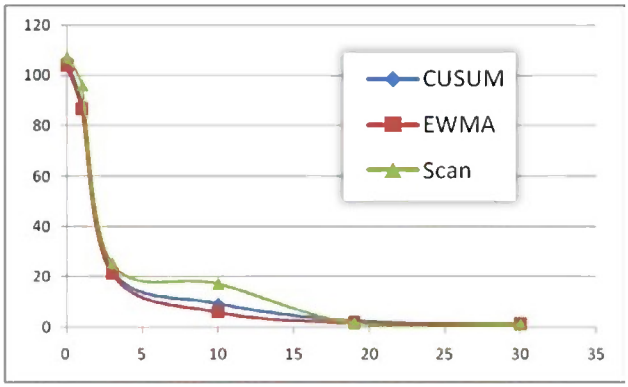

Figure 1. ADL Performance for Network Change.

5. AL-QAEDA EXAMPLE

The Center for Computational Analysis of Social and Organizational Systems (CASOS) at Carnegie Mellon University collected and maintains social network data on the Al-Qaeda terrorist network. This data includes many different relationships to include communication, financial, physical, etc. The data set begins with intelligence collected in 1988 and includes consecutive years through 2004. Using this data and SNA methods, analysts are able to calculate and quantify the most influential terrorists in the network, the most knowledgable, individuals that connect separate subsections within the group, and much more. While it is important to understand terrorist organizations from an SNA perspective, it does not necessarily identify critical changes in social network structure over time.

Social network measures were plotted over time for the number of agents, the average degree, the average betweenness, the average closeness, the average eigenvector centrality, and the density. Each of these network measures were increasing from 1988 until 1994. The measures then leveled off. There are many possible reasons for this burn-in period, the least of which is the quality of intelligence gathering on Al-Qaeda. For this reason, the average measure and standard deviation were calculated over five years beginning in 1994. The CUSUM control chart was used to monitor the five measures above from 1994 to 2001. Figure 2 displays the plot of the social network measure for the average betweenness of members in the Al-Qaeda network. It can be seen in Figure 2 that the increase in average betweenness is very slight over the time frame between 1994 and 2001. An SPC method can highlight organizational change more rapidly.

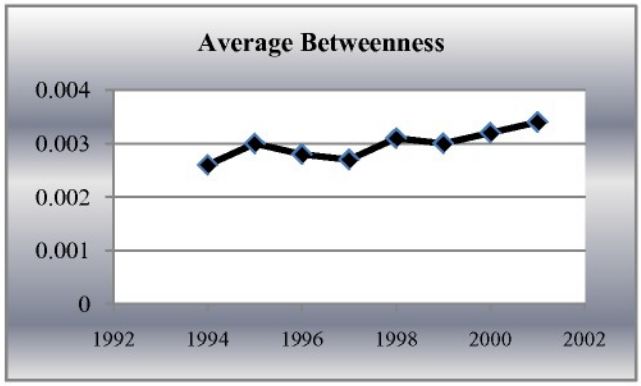

Figure 2. Average Betweenness of Al-Qaeda over time.

Figure 3 displays the plot of the CUSUM control chart statistic from 1994 to 2004. The CUSUM used a reference value of 0.5 and was calibrated to have a false alarm rate of 0.01 or once every 100 observations.

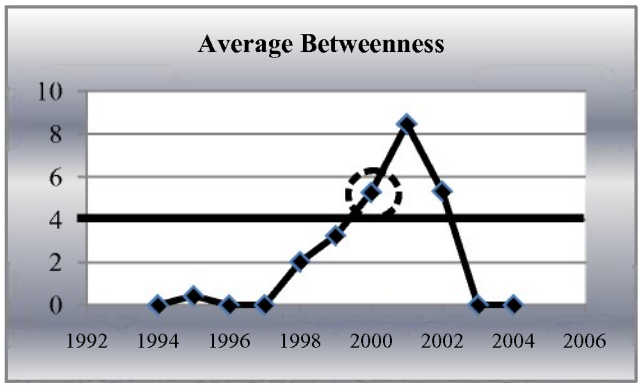

Figure 3. CUSUM Statistic for the Average Betweenness of Al-Qaeda.

It can be seen that the CUSUM statistic in Figure 3 is a more dramatic indication of network change than simply monitoring the network measure in Figure 2. This is a result of the CUSUM statistic taking into account previous observations of the network. A single observation of a network measure that is slightly higher than normal may not indicate a change in the network. Multiple observations that are slightly higher than normal, however, may indicate a shift in the mean of the measure.

Figure 3 also shows the CUSUM statistic exceeding the decision interval of 4 in 2000. This suggests that there mi exist an important change in the network structure of Al-Qaeda prior to 2000. This procedure would alert an analyst to a potential change in this terrorist organization prior to the terrorist attacks of September, 11 2001.

In order for the analyst to determine the most likely cause of the change, he would need to look at events prior

to the actual detection of the change. One method for identifying a likely change point, is to start investigating the last time that the CUSUM statistic was equal to 0. There is no similar test for the EWMA or SS, however, several methods for change point detection are available in the literature (Pignateillo and Samuel, 2001). For all network measures investigated, the change point is estimated as 1997. An analyst must therefore research events occurring in 1997 to determine the likely cause of change in the Al-Qaeda terrorist organization.

Details of Al-Qaeda's activity in 1997 are outlined in McCulloh, Carley, and Webb (2007). "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). 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 timings, they must have been working throughout 1997 planning future terrorist operations (Marquand, 2001). In February of 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. Six months later the US embassies in Tanzania and Kenya were bombed by Al-Qaeda. Essentially, 1997 was possibly the most critical year in uniting Islamic militants and organizing Al-Qaeda for offensive terrorist attacks against the United States."

6. CONCLUSION

Statistical process control is an important tool in manufacturing and may be an important tool for quickly detecting network change. The Al-Qaeda example and the evidence from multi-agent simulations demonstrate that social network monitoring could enable analysts to detect important changes in networks over time. With the CUSUM, the most likely time that the change occurred can also be determined[1]. 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.

It is important to point out that the validity of these results on real world data has not been established. The Al-Qaeda example is used to demonstrate the importance of research in this area and provide an application context. The data was based on open-source, incomplete information. It is not clear that a complete or accurate representation of Al-Qaeda is contained in the data. However, identifying this key change in the data using a statistically sound approach is encouraging. It shows the promise of predictive methods for network analysis. Further evidence of the significance of the likely change point provides even greater support of dynamic network change detection success. The simulation results are thus validated in this context and are therefore likely to provide good estimates of change detection performance.

The application of statistical process control to network science has many important applications. This method provides insight into policy decisions affecting network change. It provides a quantitative measure of system change in network structures. It also shows promise for predictive analysis of networks. This approach can also provide early warning of change, allowing senior military leaders to respond more effectively to threats and exploit success. Immediate applications to command and control, counter terrorism, and drug interdiction are obvious. As network science continues to develop, we will likely see more research in this exciting and promising field.

[1] Two social network change detection algorithms (Shewhart X-Bar and the Cumulative Sum) 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.cmu.edu.

ACKNOWLEDGEMENTS

This is a project of the U.S. Military Academy Network Science Center and the Center for Computational Analysis of Social and Organizational Systems (CASOS) at Carnegie Mellon University. More detail can be found in Army Research Institute Technical Report 1235 and Carnegie Mellon School of Computer Science, Institute for Software Research Technical Report 08-135. This work was supported in part by the Army Research Institute for the Behavioral and Social Sciences, Army Project No. 611102B74F, the Army Research Labs Grant No. DAAD19-01-2-0009 , the Office of Naval Research (ONR), United States Navy Grant No.N00014-02-10973 on Dynamic Network Analysis, the Air Force Office of Sponsored Research (MURI: Cultural Modeling of the Adversary Organization, 600322) , and the NSF IGERT program in CASOS (DGE-9972762).

REFERENCES

Baller, D., Lospinoso, J., and Johnson, A.N. (2008). An Empirical Method for the Evaluation of Dynamic Network Simulation Methods. In *Proceedings, The 2008 International Conference on Information and Knowledge Engineering, Las Vegas, NV*, 358-364.

Carley, K.M. (1990). Group stability: A socio-cognitive approach. In Lawler E., Markovsky B., Ridgeway C., and

Walker H. (Eds.) *Advances in group processes: Theory & research*. VII., 1-44. Greenwhich, CN:JAI Press.

Carley, Kathleen 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. (2003). Dynamic Network Analysis. In *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers, Ronald Breiger, Kathleen Carley, and Philippa Pattison, (Eds.) Committee on Human Factors, National Research Council*, National Research Council. Pp. 133-145.

Carley, K.M., Columbus, D., DeReno, M., Reminga, J. and Moon, I. (2008). ORA User's Guide 2008. *Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report, CMU-ISR-08-125.*

Committee on Network Science for Future Army Applications, National Research Council. (2005). *Network Science.* The National Academies Press.

Fisher, R.A., and Mackenzie, W. (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.

Freeman, L. (1979). Centrality in social networks: I, conceptual clarification. *Social Networks* 1 (1979), 215-239.

Freeman, L. (1977). A set of measures of centrality based on betweenness. *Sociometry,* **40** (1977), 35-41.

Jehl, D. (1997,). Islamic Militants Attack Tourists in Egypt. *The New York Times*, November 23, 1997. p. WK2.

Marquand, Robert (2001). The tenets of terror. *Christian Science Monitor*,18 Oct 2001.

McCulloh, I.A. (2004). *Generalized Cumulative Sum Control Charts.* Master's Thesis. Tallahassee, FL: Florida State University.

McCulloh, I.A., Carley, K.M. (n.d.). *Social Network Change Detection.* Submitted to the Journal of Social Networks.

McCulloh, I.A., Carley, K.M. (2008a). *Detecting Change in Human Social Behavior Simulation*. Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report CMU-ISR-08-135

McCulloh, I.A., Carley, K.M. (2008b). *Social Network Change Detection.* Carnegie Mellon University, School of Computer Science, Technical Report, CMU-CS-08-116.

McCulloh, I.A., Lospinoso, J., Carley, K.M. (2008). Network Simulation Models. In *Proceedings of the 26th Army Science Conference, Orlando, FL. In press.*

McCulloh, I.A., Webb, M., Carley, K.M., Graham, J. (2008). *Detecting Changes in Social Networks Using Statistical Process Control.* Army Research Institute Technical Report 1235.

McCulloh, I.A., Carley, K.M., Webb, M. (2007). *Social Network Monitoring of Al-Qaeda.* Network Science. Vol 1, Issue 1. October 2007. 25-30.

Montgomery, D.C. (1991). *Introduction to Statistical Quality Control*, 2nd Edition, John Wiley and Sons, New York.

Page, E.S. (1961). Cumulative Sum Control Charts. *Technometrics* 3, 1-9.

Pignatiello, J.J., Jr. and Samuel, T.R. (2001) Estimation of the change point of a normal process mean in SPC applications. *Journal of Quality Technology*, 33, 82-95.

Roberts, S.V. (1959) Control chart tests based on geometric moving averages. *Technometrics* 1, 239-250.

Ryan, T. P. (2000). *Statistical Methods for Quality Improvement.* 2nd Ed, Wiley.

Schreiber, C. & Carley, K. (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.

Snijders, T.A.B. (2005). Models for Longitudinal Network Data. Chapter 11 in P. Carrington, J. Scott, & S. Wasserman (Eds.), *Models and methods in social network analysis*. New York: Cambridge University Press.

Schweinberger, M., (2005). Statistical modeling of network panel data: Goodness of fit. Unpublished manuscript http://www.stat.washington.edu/msch/goodness-of-fit.pdf