Quality Engineering

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/lqen20

Research in network monitoring: Connections with SPM and new directions

Nathaniel T. Stevens, James D. Wilson, Anne R. Driscoll, Ian McCulloh, George Michailidis, Cecile Paris, Peter Parker, Kamran Paynabar, Marcus B. Perry, Mostafa Reisi-Gahrooei, Srijan Sengupta & Ross Sparks

To cite this article: Nathaniel T. Stevens, James D. Wilson, Anne R. Driscoll, Ian McCulloh, George Michailidis, Cecile Paris, Peter Parker, Kamran Paynabar, Marcus B. Perry, Mostafa Reisi-Gahrooei, Srijan Sengupta & Ross Sparks (2021) Research in network monitoring: Connections with SPM and new directions, Quality Engineering, 33:4, 736-748, DOI: 10.1080/08982112.2021.1974035

To link to this article: https://doi.org/10.1080/08982112.2021.1974035

Published online: 14 Oct 2021.

QUALITY ENGINEERING
2021, VOL. 33, NO. 4, 736–748
https://doi.org/10.1080/08982112.2021.1974035

Taylor & Francis
Taylor & Francis Group

THE PAST, PRESENT, AND FUTURE OF NETWORK MONITORING: A PANEL DISCUSSION

Research in network monitoring: Connections with SPM and new directions

Nathaniel T. Stevens[a*], James D. Wilson[b*], Anne R. Driscoll[c**], Ian McCulloh[d,e**], George Michailidis[f**], Cecile Paris[g**], Peter Parker[h**], Kamran Paynabar[i**], Marcus B. Perry[j**], Mostafa Reisi-Gahrooei[f**], Srijan Sengupta[k**], and Ross Sparks[g**]

[a]Statistics & Actuarial Science, University of Waterloo, Waterloo, Ontario, Canada; [b]University of Pittsburgh, Pittsburgh, Pennsylvania; [c]Virginia Polytechnic Institute and State University, Blacksburg, Virginia; [d]Accenture Federal Services, Washington, District of Columbia; [e]Johns Hopkins University, Baltimore, Maryland; [f]University of Florida, Gainesville, Florida; [g]Computational Informatics, CSIRO, Sydney, New South Wales, Australia; [h]Aeronautics Systems Engineering Branch, NASA Langley Research Center, NASA, Hampton, Virginia; [i]Georgia Institute of Technology, Atlanta, Georgia; [j]Information Systems, Statistics, and Management Science, University of Alabama, Tuscaloosa, Alabama; [k]Statistics, North Carolina State University, Raleigh, North Carolina

ABSTRACT
Traditional statistical process monitoring (SPM) provides a useful starting point for framing and solving network monitoring problems. In this paper the panelists discuss similarities and differences between the two fields and they describe many challenges and open problems in contemporary network monitoring research. The panelists also discuss potential outlets and avenues for disseminating such research.

KEYWORDS
graphs; network monitoring; network science; statistical process monitoring; surveillance

Question

Practitioners in quality engineering, statistics, and beyond have long been accustomed to statistical process monitoring (SPM) and change-point analyses, as a means to detect change in multivariate processes. What, in your view, sets network monitoring apart from these more traditional monitoring problems?

Reisi and Paynabar

There is a close connection between statistical process monitoring (SPM) and network monitoring. SPM is mainly concerned with detection of a change in the mean and/or covariance structure of a sequence of multivariate variables or features. A network monitoring problem can be translated into a multivariate process monitoring problem by extracting a set of features or parameters of networks to be monitored using SPM methods. For example, one can extract a group of centrality measures, including degree, betweenness, and closeness centrality, and monitor them using multivariate CUSUM or EWMA charts (McCulloh and Carley 2011; Woodall et al. 2017). More complex features, defined based on scan statistics, have also been used for network monitoring, which are more suitable for local change detection (Priebe et al. 2005). In contrast to the feature-based approaches, model-based network monitoring approaches construct a statistical or probabilistic model, such as a stochastic block or a multivariate Poisson model, that explains the interconnection among the nodes in the network. The parameters of the model are then monitored using SPM (Dong, Chen, and Wang 2020; Wilson, Stevens, and Woodall 2019; Yu, Woodall, and Tsui 2018). Statistical models such as logistic regression and generalized linear models have also been considered for modeling a sequence of attributed networks (Ebrahimi et al. 2021; Gahrooei and Paynabar 2018; Azarnoush et al. 2016). These approaches resemble the profile monitoring techniques in SPM where functional data are summarized by a set of parameters associated with basis functions that are monitored using SPM methods. In general, network modeling or feature extraction is a necessary step for network monitoring.

Another connection between SPM and network monitoring pertains to the application of network monitoring techniques to multivariate process

CONTACT Nathaniel T. Stevens nstevens@uwaterloo.ca Statistics & Actuarial Science, University of Waterloo, Waterloo, Ontario, Canada; James D. Wilson wilsonj41@upmc.edu University of Pittsburgh, Pittsburgh, Pennsylvania, USA.
*Discussion Editors.
**Discussion Panelists.
[1]https://www.cisco.com/c/en/us/solutions/automation/what-is-network-monitoring.html.

monitoring. One may view network monitoring from the covariance monitoring perspective. Network monitoring aims to detect changes in the interactions among components. If each of the variables of a multivariate process are modeled as nodes of a network, connected through their level of correlation (or a general similarity measure), then a sequence of networks is created from multivariate data streams, and can be monitored thusly.

It should be noted that although a variety of multivariate monitoring problems have been studied in the SPM literature – including change detection in autocorrelated and dynamic processes – similar problems have yet to be addressed in network monitoring. For example, temporal dependency of the network snapshot has not been adequately studied. A handful of studies integrated state-space formulation into the network modeling to capture the natural dynamics of the networks (Ebrahimi et al. 2021; Gahrooei and Paynabar 2018; Zou and Li 2017). Additionally, more work is required to study extraction of the best set of features or an appropriate model that captures both spatial and temporal characteristics of a network sequence. For example, how to capture the sparsity, the modularity, and complex heterogeneous temporal patterns of dynamic networks are the questions to be addressed for designing a successful network monitoring approach.

Sengupta

A network of n nodes is usually represented as a binary adjacency matrix, A, such that $A_{ij} = 1$ if nodes i and j are connected, and $A_{ij} = 0$ otherwise. A time series of networks can be represented as a time series of adjacency matrices, denoted by $\{A(1), ..., A(t), ...\}$. It is therefore natural for SPM practitioners to interpret each adjacency matrix as a multivariate observation and view network monitoring as an example of multivariate process monitoring. While this is a legitimate viewpoint from first principles, a few distinct aspects of network data should be kept in mind.

First, each adjacency matrix is composed of n^2 random variables for directed networks and $\binom{n}{2}$ random variables for undirected networks. Even for moderate values of n, say $n = 100$, the number of variables, 104, becomes much larger than what is typically encountered in multivariate SPM. In fact, it is common for the number of variables to be larger than the number of time points, which makes the problem high-dimensional in addition to multivariate.

In order to reduce the dimension of the problem, network monitoring is often carried out by focusing attention on a particular set of network summaries (e.g., nodal measures of centrality), rather than the entire adjacency matrix. While this seems akin to multivariate SPM, there is an important difference. In multivariate SPM, typically the variables being monitored are well-defined and directly observed. In contrast, network summaries used for monitoring are being defined by the user as functions of the observed network data, and may not be complete descriptions of network behavior.

Another important distinction concerns the inherent interdependency of interactions within a network. Such interdependency structures tend to be more complex than what is typically encountered in multivariate SPM. For example, if the behavior of a node changes over time, this will affect its neighbors, and it's neighbors' neighbors, and those neighbors' neighbors, and ultimately the entire network at the same time. Accounting for such dependency is an important consideration in network monitoring methodology.

Driscoll

To answer this question in one word: complexity! Networks are complex, and more complex than other data structures encountered in traditional SPM applications. When applying traditional process monitoring techniques to industrial settings, in many cases, the statistic being monitored is straightforward to measure. We have methods that work very well for data measured on a continuous scale, and we have also developed techniques that handle the complexities associated with attribute-type data. We have also developed methods for more complex multivariate data. In each case, we also have an understanding of the benefits and shortcomings of the available methodology. However, network data tends to combine many of the individual complexities seen in industrial applications into a single dataset.

For instance, multivariate issues arise when considering a time series of observed matrices, and we must also acknowledge that networks are expected to evolve over time. Thus, methods (such as scan-based ones) that accommodate gradual evolution as well as seasonal changes must be developed. Furthermore, we have an arsenal of tools to model network data, but as these models become more flexible and more realistic, they become increasingly complex, and so too do the associated monitoring methods.

Another complexity is the sheer size of network data. We are well aware that we live in an era of big data where data is being produced and collected at staggering rates. This is especially true with network data. I agree with many current researchers in network monitoring who advocate for the use of simulation to study the properties of proposed monitoring techniques. Depending on the context of the network monitoring problem, the computational burden of such investigations can be very high. Some simulations can take weeks or even months to run. We see this issue with some applications in industrial settings, but the prevalence and magnitude of this problem are both larger in network monitoring contexts.

Michailidis

There are two flavors of network monitoring, the offline (retrospective) and the online (prospective) ones. The former assumes that one has collected univariate/multivariate data evolving over time and the goal becomes to identify if there exist change-points and what their locations are, retrospectively, assuming some model/mechanism for their temporal evolution. On the other hand, the online version assumes that new data are obtained in real time and one is interested in identifying a change from the current regime that governs their evolution. There is a very mature body of literature for both flavors for univariate and multivariate problems and for various types of statistical models employed to formulate the detection problem. These include parametric models for changes in the mean or the variance/covariance of the data generating distribution, models for changes in the temporal dynamics of the process, as well as nonparametric approaches.

However, network models offer a number of novel challenges, including issues related to size, the nature of the change, and also the type of network model assumed. A network consisting of n nodes can have up to n^2 edges, in the case of a fully connected network. It can then easily be seen that network monitoring problems become high dimensional in nature, especially for processes related to the edges. This presents both computational and technical challenges. Note that in traditional change-point analysis techniques involving a few hundred processes, the computational bottleneck comes primarily from the time dimension T. Specifically, the goal is to come up with algorithms with near linear complexity in T and avoid dynamic programming-based ones that exhibit quadratic complexity, despite their generality and (near) optimal detection properties. However, in network problems the node and consequently the edge dimension can be much larger than the time dimension ($n \gg T$). Hence, the former becomes the leading factor in determining the computational complexity of detection algorithms. Another complication arising from the large network dimension is that the data may be stored across various locations, whose processing may impose additional computational burden.

The computational complexity due to the network dimension is also related to the nature of the change. For example, suppose that the network change is "global" in nature, in the sense that a large portion of the nodes/edges is impacted. In that case, it is reasonable to devise a detection algorithm that samples a subset of nodes/edges and tests for the presence/absence, or estimates the location, of the underlying change-points. On the other hand, if the change is "local" in nature, in the sense that only a small set of nodes/edges are impacted, then the problem becomes significantly harder. The total available signal becomes fairly small, which becomes harder to detect, and usually requires the construction of more refined detection algorithms that are computationally expensive. This problem can be mitigated if domain knowledge can aid in selecting small sets of nodes or edges that are of primary interest and only monitor those, or the network topology has a known a priori structure that can be leveraged accordingly. An example for the latter setting may correspond to a network with community structure, where it is known that changes are primarily driven by creation or deletion of connections (edges) within the community.

Finally, from a technical viewpoint, network structures can introduce intricate dependencies amongst the random variables that define the corresponding change-point detection problem. Hence, obtaining rigorous results on the performance properties of detection algorithms becomes significantly more challenging.

McCulloh

The key difference in network monitoring compared with SPM is the network treatment of highly dependent data, where IID assumptions do not hold. Network methods essentially offer methods to transform highly dependent relational data into a graph structure where measures can be treated much like traditional SPM methods.

The fields of social network analysis and network science are also able to apply a useful body of knowledge regarding the behavior of certain types of

networks that can inform analysis, reduce variance, and result in more statistically powerful analysis. For example, in a typical social network of random interactions, centrality measures of degree, closeness, betweenness, and eigenvector centrality are highly correlated. As the network exhibits more and more structure, anomalies occur where a node may take on a role as a boundary spanner between two or more subgroups. In this case, their degree may be relatively low when compared to betweenness. It is often these tradeoffs that are most interesting and provide the most relevant insights. Thus, the application of multivariate SPM is highly relevant and useful.

Whether an engineer applies SPM methods to graph-level data, multivariate measures from the network, or some other transformation, network methods offer many different ways to treat relational data. Furthermore, the combinatorics of different combinations of dependent data increases the complexity of the sample space. This complexity opens up exponentially greater applications and scenarios to explore from a scientific perspective.

To compound the complexity are the simultaneous rise in network-based time series methods. Stochastic Actor-Oriented Models (SAOM) also known as Siena models provide a statistical framework for co-evolving nodal and relational variables and their interactions over time (Snijders, Van de Bunt, and Steglich 2010). Temporal Exponential Random Graph Models (TERGM) provide a similar approach where network structure at one time point may inform future structure (Hanneke, Fu, and Xing 2010). These frameworks introduce new methods to analyze and treat networks over time that may both compete with statistical network monitoring as well as provide new models to consider.

Perry

There are many notable differences between statistical network monitoring problems and more traditional statistical process monitoring (SPM) problems. First, in traditional SPM problems, an assumption is often made about the parametric distribution of the charting statistic, usually based on distributional assumptions of the data. Consequently, parametric modeling strategies are quite commonplace. However, in network monitoring problems, the distribution of the charting statistic is rarely (if ever) known, which implies perhaps nonparametric or model-free monitoring strategies should be considered. In traditional distribution-free SPM methods, although they require no distributional assumptions, most require the monitored data sequence to be independent, and hence free of autocorrelation. Unfortunately, this assumption leads to tools with very limited impact on today's modern quality engineering practice. Recently, however, distribution-free monitoring strategies have been developed that consider autocorrelation and may have direct application to network monitoring problems. For example, Perry and Wang (2020) developed a distribution-free joint monitoring strategy for ergodic-stationary univariate continuous processes, which can permit effective monitoring of almost any covariance-stationary network charting statistic while providing adequate control of the false alarm rate. Other plausible methods are given in Qiu, Li, and Li (2020) for univariate processes and Xue and Qiu (2021) for multivariate processes. In the same sense, it is also my view that categorical process monitoring methods can play a significant role in the future of statistical network monitoring. For example, Perry (2020) proposed a classification rule based upon the observed values of two different network statistics (one that measures reciprocity and another that measures transitivity), in an effort to classify a network as being hierarchical or not. Consequently, their network monitoring problem was reduced to monitoring a simpler categorical process, where the false alarm rate was easier to control and the interpretation of a control chart signal was relatively straightforward.

Autocorrelation is not the only type of correlation to be considered when monitoring networks. Homophily, for example, also induces correlation. With this form of dependency, the likelihood of any two vertices being connected depends on the value of one or more attributes measured at the vertices, e.g., vertices with similar attributes are more likely to be connected. Similarly, network motifs are correlated. Consider the simple case of monitoring the number of triangles in a Bernoulli graph over time, where the edges between vertices are IID Bernoulli random variables with parameter p. One might be tempted to model the triangle count as a binomial random variable and then apply a binomial-based control chart for monitoring. However, this would be flawed since each triangle in the network shares an edge with $3n - 9$ other triangles, where n is the number of network vertices. Specifically, Perry (2019) showed that the covariance between any two triangles, say X and Y, in such a network is given by

$$\mathrm{Cov}(X, Y) = \begin{cases} p^3(1-p^3) & \text{if } X \text{ and } Y \text{ are same triangle} \\ p^5(1-p) & \text{if } X \text{ and } Y \text{ share an edge} \\ 0 & \text{if } X \text{ and } Y \text{ do not share an edge} \end{cases}$$

and consequently, although the edge counts follow a binomial distribution, the triangle counts cannot follow a binomial distribution since triangle pairs are not necessarily independent.

One additional, and very important, difference between network monitoring and more traditional SPM applications is the fact that many aspects of networks are continually changing, and oftentimes the observed change is not an anomalous one. This poses additional challenges for the statistical process monitoring practitioner as network charting statistics should be carefully chosen such that they sufficiently capture the precise network change phenomenon of interest.

Sparks and Paris

The multivariate monitoring of time between communication events in social media networks is difficult because events do not generally occur simultaneously, and this makes it difficult to use time between events for this purpose. If we are interested in multivariate SPM, then counting processes may be the way to go. However, the focus would be on counts that are expected to correlate quite well. The inherent multivariate structure of the social network, particularly with direct network connections within very large networks makes this a very challenging task. A potential solution is to decompose a large network into disjoint sub-networks if this is possible without the loss of too much information. This may make the task more manageable for very large social networks.

Another challenge associated with social media monitoring that is typically not seen in traditional SPM applications is the development of robust practical plans that can be used in applications that are efficient at flagging outbreaks for a range of sizes and starting points in steady state situations. The zero state situations are less common in social networks, and are more aligned with the manufacturing of parts or items.

Parker

As someone who does not actively participate in network monitoring (NM) research, but who has a strong background in quality engineering and traditional SPM applications, the discussion's editors invited me to provide my outside perspective. Having conducted a broad, high-level literature survey on the topic, I found both amiable similarities and striking differences between these monitoring endeavors. As in most cases of emerging research applications, what's new is often built on and leverages what's old; that's if the research is done well. Of course, that's not meant to imply that existing methods are sufficient in tackling the new challenges of NM. However, I believe that the conceptual and philosophical foundation of existing quality methods offer a beneficial perspective for NM.

Clearly, in today's society, networks are ubiquitous and serve scientific, health, governmental, industrial, and personal functions. Statistical research in network monitoring has applications in a variety of diverse domains where it is important to characterize, monitor, and detect anomalies in the dynamics between interconnected entities that exchange information within a network. To manage the scope of my discussion, I envisioned a physical computer network as a tangible, motivating case study. However, I believe that my observations and suggestions are easily mapped to other network domains.

Definitions are always a good place to start to make sure we're solving the "right" problem. When I think about maximizing research impact, I tend to favor definitions coming from the stakeholders, the users, the beneficiaries of the statistical research. Consistent with that approach and considering physical computer networks as a motivating case study, I found a particularly helpful answer to the question "What is Network Monitoring?" from Cisco[1]: "*Network monitoring provides the information that network administrators need to determine, in real time, whether a network is running optimally. With tools such as networking monitoring software, administrators can proactively identify deficiencies, optimize efficiency, and more.*" This definition is broadly applicable to many network domains, and I will endeavor to highlight a few of its salient aspects to frame my remarks.

In the first sentence, the phrase "provides the information" highlights that NM provides actionable information to make decisions on the state of a network. Furthermore, it is assumed that information will be partial or incomplete in some sense, and therefore the subsequent decisions derived from it will have inherent uncertainty. Supporting data-driven decisions made in the presence of uncertainty drives research into statistical methods, which confirms that this is a backyard, as Tukey would say, where statisticians could and should play.

Next, the definition points us to the specific decision context supported by an ultimate utility of research in statistical methods for NM, namely "to determine, in real time, whether a network is running optimally." An important feature of this phrase is

"real time", which seems to inevitably suggest that the rich, mature library of statistical process monitoring approaches should offer insights on how to approach this new challenge. Obviously, the difference from traditional SPM for computer networks, is that we are monitoring patterns of bits and bytes, not physical widgets coming off a factory assembly line. However, whether the patterns are digital signatures or parameterized features of a physical part, both are mathematical abstractions that allow us to differentiate common cause and special causes of variation.

To detect a special cause of variation, a cause for an alarm, both NM and SPM require learning and describing what's normal. In SPM, this is usually from a retrospective analysis of historical data, known as Phase I analysis, that considers the system's random variability, and derives boundaries that indicate whether our system is under normal operations, or if it has departed from normal. An important assumption in a Phase I analysis is that the historical data are largely representative of the data to be monitored in the future. However, networks are dynamic, changing and evolving over time, and so a Phase I assumption of stationarity may not be fully applicable. In industrial SPM, there was a branch of applications in actively adapting processes that used feedback controllers; they may offer helpful ideas when actively monitoring a dynamic system. See Jiang et al. (2002) for an introduction to this application area.

Interestingly enough, the potential causes of deviation from normal operations in NM are often quite different than in SPM, since they include the potential for adversarial, intentional disruptions. It's not common, or at least not commonly published in the literature, for a traditional SPM method to be designed to detect industrial sabotage, however this certainly would be a special cause. In contrast to quickly detecting system degradation of production machinery, adversarial disruptions could come without warning and present themselves as more of a spike in behavior and might disappear immediately afterwards. Once again, in contrast to traditional SPM, we often don't consider the situation fixing itself once a signal occurs. Admittedly, I'm generalizing both NM and traditional SPM applications, however my main point is that it's important to recognize the potential sources of a disruption to strategically design a system to effectively detect them, and I believe the types of disruptive signatures differ between NM and traditional SPM.

Another aspect of Cisco's definition is "to determine … if a network is running optimally". Frequently, when consulting with subject-matter experts, they often say they want to "optimize" their system. When I ask them to define the "optimal" state quantitatively, I usually get a response that would be better described as wanting to "improve" their system, i.e., find a trajectory of improvement, rather than seeking some utopian, optimal state where they would cease all further research. In practice, the optimal state is often practically unattainable, since it depends on many competing factors. I share this difference between the objective to "improve" and "optimize," in order to highlight that NM monitoring seeks continuous improvement by identifying factors and their relationships with responses that indicate that the dynamic, evolving system is moving in a direction of "goodness." While this may sound like steepest ascent in response surface methodology would be directly applicable, my previous comments on lack of system stationarity combined with a high dimensional factor space may not make those techniques readily applicable. However, I could envision NM benefiting from an approach known as evolutionary operation (EVOP), proposed by Box (1957), where small perturbations are intentionally induced during normal operations to identify a path of process improvement. Developing methods of in-process designed experimentation with small, not disruptive, changes would seem to be a very applicable strategy for NM. Hahn and Dershowitz (1974) and Box and Draper (1969) are recommended for additional information on EVOP.

The second sentence in Cisco's definition indicates that there is a NM system, consisting of " … tools such as networking monitoring software, administrators can proactively identify deficiencies, optimize efficiency, and more." As with traditional SPM, it requires more than simply implementing a control chart to practically impact an organization's overall quality process, it requires a system that includes appropriate tools and techniques, decision metrics to assess goodness, and a response plan to disruptions, known as an out-of-control action plan (OCAP). It also requires trained individuals who know how to use the tools, experts who know how the tools work, and leaders who understand and appreciate the system's function and value. As a caution to researchers in the field of NM, it would be instructive to consider how SPM was eventually adopted in industry and appreciate and learn from its cultural struggles. With this emerging field of NM, I expect that there will be similar struggles, particularly since I expect the "network administrators" in most cases are not

statisticians and they may not fully appreciate the benefits of employing statistical inference. I believe that for NM techniques to have "big impact", it will likely take cultural change to adopt the statistical research being conducted.

Considering the multiple disciplines engaged in NM research, a statistical thinking approach is a unique perspective to offer. Statistical thinking discussed by Hoerl and Snee (2020) recognizes that all productive work occurs in interconnected processes that have inherent variation, and understanding and reducing the variation is a key to successful quality improvement. It appears abundantly clear that this "old" concept remains quite relevant in the emerging, maturing field of NM.

Editors' comments

The panelists agree that statistical process monitoring (SPM) methods provide a useful framework for online network monitoring, but that without extensions, modifications, and adaptations, existing SPM methods do not sufficiently address the various complexities associated with network monitoring. We agree with all of the points made by the panelists; here we simply synthesize and summarize these comments, amplifying those we feel are particularly important.

First and foremost, network monitoring problems differ relative to traditional SPM applications in that the statistic or property that should be monitored is not always clearcut. Whereas in traditional SPM applications, the monitored quantity is typically some well-defined and obvious quality characteristic, in networking monitoring applications, *what* to monitor is ambiguous. A second important distinction between networking monitoring and traditional SPM applications is the dimension of the problem. As several panelists discussed, this has implications for data storage, computational complexity, effective discrimination of signal from noise, as well as the feasibility of traditional methods when the dimension of the monitoring statistic exceeds the number of observed time points. Another important characteristic of network monitoring not easily handled by traditional SPM methods is the complex dependence that exists within and between temporal snapshots of the network. In particular, the intricate dependencies inherent to an interconnected system must be adequately accounted for. Furthermore, networks tend to be dynamic, evolving gradually and organically over time. It is therefore important to also accommodate such temporal dependence and natural change. Traditional SPM methods must also be adapted to account for a different type of change. As Dr. Parker discusses, networking monitoring methods must be able to quickly and accurately detect fleeting (as opposed sustained) changes that may possibly be a result of intentional, malicious actions.

While it is clear that SPM provides a useful starting point for framing and solving network monitoring problems, Drs. Reisi and Paynabar highlight a synergistic reciprocity whereby network monitoring methods may also be useful in framing and solving some traditional multivariate SPM problems.

Question

What opportunities and open problems do you see for graduate students and early-career academics who are interested in network monitoring research? In your experience, what journals or conference proceedings have been particularly interested in publishing research in this area?

Driscoll

My best recommendation for graduate students or early-career academics who are interested in network monitoring research is to first become proficient in some coding-based software like Python, R, etc. You learn so much more about the complexities of modeling and monitoring network data if you are the person who is writing the simulation code. Most articles that are published mention that the author is happy to share code. Request code and teach yourself how to generate data using different network models like the degree corrected stochastic block model (Karrer and Newman 2011), for instance. Most colleges and universities have servers available for computing. Learn about and use these resources, and develop best practices to make simulations run faster. Most network monitoring research requires quite a bit of computing power. I see this becoming even more important as we study more complex network data. Having these computing skills will make you invaluable to any network monitoring research group.

In terms of open problems, I would always encourage all researchers to start with the simplest problems first. For instance, begin with simpler, as opposed to complicated, network models. Over the past couple of years, researchers have begun investigating the impacts of changing aspects of the network model and the effects of these changes on monitoring performance. We should continue on this path using

comparative studies. These studies use simulation to compare different monitoring methods. A great resource for comparative studies using simulations are the papers written by Yu et al. (2021) and Kodali et al. (2020).

One particular research area that shows promise is the development of methods that not only signal when something unusual has occurred, but also have the ability to identify the key players that caused the unusual events. In Zhao et al. (2018), the researchers use clustering algorithms to identify anomalous subnetworks. It would be interesting to develop alternative approaches for the identification of change-inducing subgroups.

Another promising avenue for future research is to study false discovery rate properties in the context of network monitoring. Many of the available network monitoring techniques use multiple control charts to simultaneously to monitor the same data, so there is some concern about having a high number of false alarms. This research area is also suggested in Woodall et al. (2017).

In terms of publishing, it seems like many statistics journals are very interested in network monitoring and network analysis research. However, I would encourage researchers in our field to also try to publish in journals from other concentration areas so that we can have a broader impact with respect to this interdisciplinary problem.

Sparks and Paris

This is a fairly new field of research that is enhanced by most processes becoming digitally oriented, and therefore a Ph.D. now will set the trends in this field. Network monitoring will open-up several new fields of applications, particularly when the methodology is made available to users. For example, research is needed at the interface of psychology, classical sociology, computer science, natural language processing and statistics for research platforms such as the Vent platform (Malko et al. 2021).

Two research projects in which I have taken part, which deserve more attention, include the study of mental health and well-being, as well as monitoring how diseases spread at major events. In mental health, it is important to establish the reasons for emotional changes because social networks are dynamic in nature, as people either fall into or out of relationships, and as a consequence, their emotional state could change. In our previous work in disease spread, we attempted to establish the contribution of major sporting events in the Australian Football League to the spread of infectious diseases, and unfortunately we could not find the link between Winter-related diseases and events that may have contributed to the spread of these diseases. Our conclusion was that most Australians stayed away from major events when they were sick. The spread of diseases on major bus routes to and from work could be a potential future area of research; however getting the appropriate data could be a challenge given the privacy restrictions on collecting individual's data (Weiss 2009).

Once a network monitoring methodology is made available, potential applications will expand beyond the current horizons. At this point, I believe that conferences will welcome papers that apply efficient techniques for monitoring processes.

Reisi and Paynabar

Computational scalability for network modeling and monitoring both in terms of network size and temporal resolution is one of the major challenges in this area that requires further research. Existing literature in network monitoring is mainly focused on small to medium size networks with limited solutions for monitoring large networks such as transportation networks or social media networks. Many existing approaches are not efficient when the network size is on order of millions of nodes. For example, approaches that rely on the EM algorithm or Kalman filtering are inefficient in estimating the parameters of huge networks (Gahrooei and Paynabar 2018; Zou and Li 2017). This inefficiency translates into slow detection of changes, especially when temporal resolution is high and fast detection is essential.

The importance of computational scalability is more apparent in applications like the industrial internet of things (IIOT), in which many sensors and devices are communicating at high rates and quick detection of faults or intrusions are critical. If network monitoring techniques are to be applied to these systems, fast algorithms that can monitor and rapidly detect changes in a large-scale networked system are necessary.

Another major challenge in the area of network monitoring is the detection of small local changes in a network stream. The current literature mainly considers the problem of detecting global changes by summarizing networks into a set of global parameters (Ebrahimi et al. 2021; Dong, Chen, and Wang 2020; Wilson, Stevens, and Woodall 2019; Gahrooei and Paynabar 2018; Zou and Li 2017; Azarnoush et al. 2016). Designing algorithms that are sensitive to small

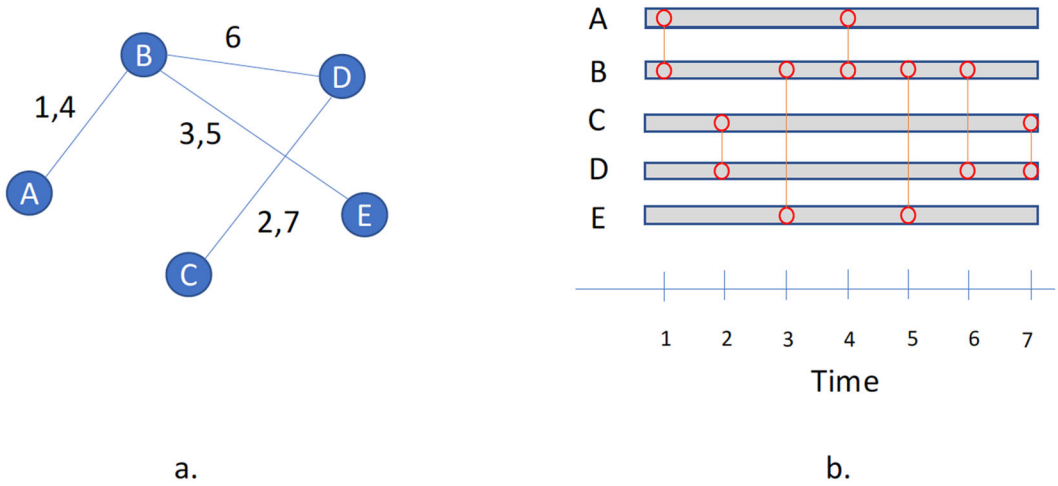

Figure 1. Illustrative example of (a) a contact graph and (b) the corresponding contact sequence.

local changes within the networks is crucial for broadening the applicability of this field to other areas such as transportation systems and cybersecurity. The main problem is that an attack to a cyber-physical system may be designed in a way that it manifests in a local and distant area of the network (Li et al. 2021). Monitoring multiplex/multilayer networks with heterogeneous interactions is another research opportunity. Many modern systems contain multiple layers of interactions including physical, correlation, and cyber interactions and should be modeled by multilayer networks (Brunetti et al. 2019). A shock or an intrusion to any layer can propagate to other layers and cause significant consequences. Therefore considering the between-layers interactions is critical for fast detection and isolating of faults (Dong, Chen, and Wang 2020). Monitoring of evolving networks is another promising topic that has not been broadly studied in the literature of network monitoring. The structure of such dynamic networks evolves over time by adding or removing edges and nodes. Many models such as scale-fee networks (preferential attachment) have been developed to study evolving networks (Barabási 2013). Nevertheless, monitoring these networks to distinguish between the removals/additions that are according to the natural evolution of the network from those that are excessive or unexpected is challenging, and has not been widely studied through the lens of SPM.

A few journals are engaged in publishing articles related to network monitoring problems. The *Journal of Quality Technology* (see for example, Gahrooei and Paynabar [2018]; Azarnoush et al. [2016]), *IISE Transactions* (e.g., Ebrahimi et al. [2021]; Woodall et al. [2017]; Zou and Li [2017]), *Quality and Reliability Engineering International* (e.g., Wilson, Stevens, and Woodall [2019]; Kan and Yang [2017]; Yu et al. [2021]), and *Technometrics* (Dong, Chen, and Wang 2020) are among the ones that are related to the quality control and SPM community. The ACM international conference on Knowledge Discovery and Data Mining has also been involved in this field (e.g., [Akoglu, McGlohon, and Faloutsos [2010]; Noble and Cook [2003]).

Perry

I believe there are several opportunities in statistical network monitoring for graduate students and early-career academics. Below I list those areas I find interesting and potentially fruitful.

Consensus monitoring: Similar to consensus clustering in complex networks, e.g., see Tandon et al. (2019), one might also consider consensus monitoring. That is, explore employing several different monitoring mechanisms and use these to form a majority consensus on whether or not critical aspects of a network have changed.

Monitoring contact sequences: By monitoring one or more contact sequences of a network, one can potentially develop strategies for detecting changes in the spread of information or disease over time. For example, Figure 1a shows a contact graph where the times of contacts between vertices are shown on the edges. Figure 1b shows a time series plot of the contacts explicitly. See Holme (2005) and Holme and Saramäki (2012) for information on contact sequences and temporal networks in general.

Monitoring network motifs: It is well known in the network science literature that higher-order local structural network configurations, i.e., motifs, are essential to understanding the fundamental structures that govern the behavior of many complex networks. For example, Holland and Leinhardt (1971) suggested transitive triads (Figure 2a) are important for

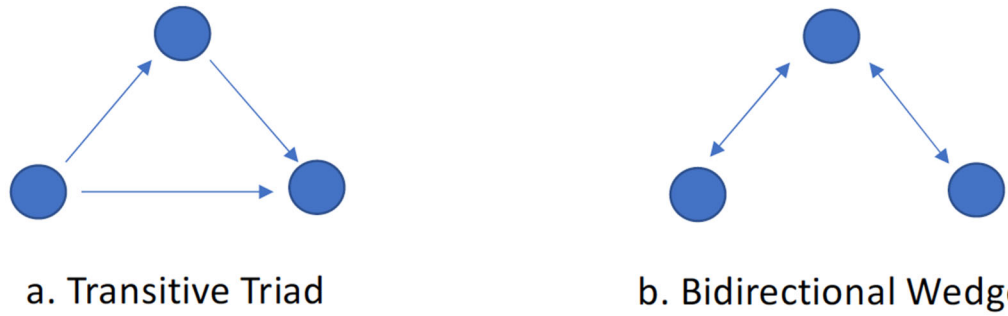

Figure 2. Two examples of network motifs.

understanding hierarchies and cliques in social networks, while Honey et al. (2007) suggested that open bidirectional wedges (Figure 2b) are important to understanding structural hubs in the human brain, where such hubs permit increased levels of information flow between distant nodes. These examples, and others, are given in Benson, Gleich, and Leskovec (2016). Consequently, monitoring for changes in the frequency of one or more such motifs of interest can potentially provide a relatively simple, yet effective means to detect interpretable change in networks.

Computational challenges: The emergence of online network-based data has led to the need for fast and scalable computational algorithms for general information extraction and processing. For example, Perry, Michaelson, and Ballard (2013) proposed a network clustering model that seeks to optimize a log-likelihood ratio statistic, often permitting the discovery of densely connected subgraphs or communities. However, computational inefficiencies only permitted clustering of relatively small networks (say <500 vertices), significantly limiting its use in practice. To overcome this limitation, Ballard and Perry (2019) proposed efficient updating equations for the model parameter estimates used in their optimization algorithm. These new contributions made the clustering strategy of Perry, Michaelson, and Ballard (2013) tractable for larger networks, and thus, much more applicable in practice. As another example, consider the task of counting the number of motifs of a given type in an observed network. Counting the occurrences of a given configuration of interest can be computationally expensive, particularly for large networks. For instance, consider counting the number of triangles in an undirected binary network. If there are n vertices, then the number of triangles can be computed from

$$T = \sum_{i=1}^{n-2} \sum_{j=i+1}^{n-1} \sum_{k=j+1}^{n} A_{ij}A_{ik}A_{jk}$$

where, in general, A_{st} denotes the $(s,t)^{\text{th}}$ element of the observed adjacency matrix A of the network. Note that the time complexity for this approach to calculating the number of triangles is $O(n^3)$, where n denotes the number of vertices in the network, and consequently, it becomes less tractable for large networks to be processed online. Furthermore, if one is to consider a more complex network motif (say, a configuration involving four vertices), then one should expect the time complexity required to count its number of occurrences to also increase. These examples both demonstrate how computational complexity can prohibit the practical use of a method, and that any new statistical network monitoring strategies should be developed with this in mind.

In my experience, there are several journals that have been open to publishing my research on networks, including *Journal of Quality Technology*, *Journal of Applied Statistics*, *Stat*, and *Computational Statistics and Data Analysis*. Additional journals where I have seen network monitoring-based research published include *Quality and Reliability Engineering International* and *IISE Transactions*.

McCulloh

Research at the intersection of SAOM and TERGM methods is particularly interesting. This would combine multiple statistical frameworks and treatments and provide trade space analysis of which methods to use under different conditions. Another area would involve the use of reduced graphs, where networks are efficiently clustered and clusters are treated as nodes in a "reduced network". This is also known as a block model, where relationships are defined between clusters. How do changes in a network compare with its reduced network? Do insights in one inform the other? Meta networks involve nodes of different types, such as people, resources, or locations. How do meta-networks impact change, changes over time, and the treatment of network data?

I have not found any journals that have focused on network monitoring and its applications. Engineering

journals tend to be disinterested in network applications. Social network journals tend to be disinterested in statistical process monitoring. Hopefully, this series of discussion papers will raise interest and increase research in this area, resulting in more scientific publications.

Michailidis

Network monitoring problems are more challenging than their counterparts for multivariate processes, due to intricate dependencies introduced by the topology and their large scale nature. Nevertheless, they offer ample opportunities for developing novel methods for change-point analysis, especially focusing on fast algorithms. Given the sheer size of networks (n nodes) and on many occasions the large number of observations (T), even linear complexity algorithms in n and T may be too expensive. Hence, there is a strong need for sublinear algorithms, possibly combined with recent advancements in sketching techniques that reduce the dimensionality of the underlying data. Another challenge comes from the fact that the available data are stored across many locations and hence distributed algorithms are required. Further, in many online monitoring applications, new data may be obtained in an asynchronous manner and thus the detection algorithms need to be able to consume such partial information. An example of the latter setting comes from sensor networks, wherein due to energy constraints, not all the sensors are always on. A similar challenge, but for different reasons arises in network surveillance, since the (possibly covert) surveyor may be able to access only part of the network.

Given the diverse applications of network monitoring, all corresponding subject matter journals are in principle interested in work developing methods tailored to problems and data in the corresponding field. More theoretically focused journals and conference proceedings in statistics, econometrics, machine learning, computer science and signal processing routinely publish work on computationally fast change-point detection strategies and their theoretical guarantees.

Sengupta

In my opinion, an important class of open methodological problems is to develop statistically rigorous network monitoring methods that allow flexible dependence between successive temporal observations. Much of the existing literature assumes that the temporal network observations are independent of each other. Accommodating dependence makes the problem significantly more challenging, but enables monitoring in a more flexible and realistic setting. I would also like to encourage greater emphasis on interdisciplinary research in network monitoring in collaboration with domain experts from various disciplines.

Generally speaking, I feel that the Quality Control and Industrial Engineering disciplines have been welcoming of network monitoring research. This includes journals such as *Quality Engineering*, *IISE Transactions*, and *Applied Stochastic Models in Business and Industry*, and conferences such as the Quality and Productivity Research Conference.

Editors' comments

The responses to the first question in this paper emphasized a number of important differences between network monitoring and traditional SPM applications. As reiterated by the panelists in this section, many of those distinctions present practical challenges and therefore interesting research problems. In particular, the development of scalable methods that may be used to monitor very large networks (in a timely manner and while accounting for the multiple comparison problem) is of interest. The development of methods that flexibly accommodate arbitrary correlation structure within and between temporal snapshots of the network is similarly of interest. Comparative studies designed to evaluate the practical and theoretical performance of existing methods are also important.

With respect to dissemination of research, like many of the panelists, we have found the quality engineering community to be receptive to network monitoring research. However, like Dr. McCulloh, we have personally found other communities to be much less welcoming. Although network science is a very multidisciplinary field, these disciplines seem to be siloed and disinterested in each other's literature and perspectives. Like Dr. McCulloh, we also hope that this series of discussion papers encourages the continued growth of network monitoring research, and we hope that it begins to break down existing interdisciplinary barriers, ultimately fostering greater innovation and impact in this area.

ORCID

Nathaniel T. Stevens http://orcid.org/0000-0001-6149-5797
James D. Wilson http://orcid.org/0000-0002-2354-935X
Ian McCulloh http://orcid.org/0000-0003-2916-3914

George Michailidis http://orcid.org/0000-0002-3676-1739
Kamran Paynabar http://orcid.org/0000-0002-6906-3611
Srijan Sengupta http://orcid.org/0000-0001-6889-8599
Ross Sparks http://orcid.org/0000-0001-5852-5334

References

Akoglu, L., M. McGlohon, and C. Faloutsos. 2010. Oddball: Spotting anomalies in weighted graphs. In Pacific-Asia Conference on Knowledge Discovery and Data Mining, Hyderabad, India, 410–421. Springer.

Azarnoush, B., K. Paynabar, J. Bekki, and G. Runger. 2016. Monitoring temporal homogeneity in attributed network streams. *Journal of Quality Technology* 48 (1):28–43. doi:10.1080/00224065.2016.11918149.

Ballard, A., and M. B. Perry. 2019. Improving computational performance in likelihood-based network clustering. *Stat* 8 (1):e256. doi:10.1002/sta4.256.

Barabási, A.-L. 2013. Network science. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 371 (1987):20120375. doi:10.1098/rsta.2012.0375.

Benson, A. R., D. F. Gleich, and J. Leskovec. 2016. Higher-order organization of complex networks. *Science (New York, N.Y.)* 353 (6295):163–6. doi:10.1126/science.aad9029.

Box, G. E. 1957. Evolutionary operation: A method for increasing industrial productivity. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 6 (2): 81–101.

Box, G. E., and N. R. Draper. 1969. *Evolutionary operation: A statistical method for process improvement*. New York City, NY: Wiley.

Brunetti, C., J. H. Harris, S. Mankad, and G. Michailidis. 2019. Interconnectedness in the interbank market. *Journal of Financial Economics* 133 (2):520–38. doi:10.1016/j.jfineco.2019.02.006.

Dong, H., N. Chen, and K. Wang. 2020. Modeling and change detection for count-weighted multilayer networks. *Technometrics* 62 (2):184–95. doi:10.1080/00401706.2019.1625812.

Ebrahimi, S., M. Reisi-Gahrooei, K. Paynabar, and S. Mankad. 2021. Monitoring sparse and attributed networks with online hurdle models. *IISE Transactions* :1–14. doi:10.1080/24725854.2020.1861390.

Gahrooei, M. R., and K. Paynabar. 2018. Change detection in a dynamic stream of attributed networks. *Journal of Quality Technology* 50 (4):418–30. doi:10.1080/00224065.2018.1507558.

Hahn, G. J., and A. F. Dershowitz. 1974. Evolutionary operation to-day—some survey results and observations. *Applied Statistics* 23 (2):214–8. doi:10.2307/2347004.

Hanneke, S., W. Fu, and E. P. Xing. 2010. Discrete temporal models of social networks. *Electronic Journal of Statistics* 4:585–605. doi:10.1214/09-EJS548.

Hoerl, R. W., and R. D. Snee. 2020. *Statistical thinking: Improving business performance*. Hoboken, NJ: John Wiley & Sons.

Holland, P. W., and S. Leinhardt. 1971. Transitivity in structural models of small groups. *Comparative Group Studies* 2 (2):107–24. doi:10.1177/104649647100200201.

Holme, P. 2005. Network reachability of real-world contact sequences. *Physical Review. E, Statistical, Nonlinear, and Soft Matter Physics* 71 (4 Pt 2):046119. doi:10.1103/PhysRevE.71.046119.

Holme, P., and J. Saramäki. 2012. Temporal networks. *Physics Reports* 519 (3):97–125. doi:10.1016/j.physrep.2012.03.001.

Honey, C. J., R. Kötter, M. Breakspear, and O. Sporns. 2007. Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *Proceedings of the National Academy of Sciences of the United States of America* 104 (24):10240–5. doi:10.1073/pnas.0701519104.

Jiang, W., H. Wu, F. Tsung, V. N. Nair, and K. L. Tsui. 2002. Proportional integral derivative charts for process monitoring. *Technometrics* 44 (3):205–14. doi:10.1198/004017002188618392.

Kan, C., and H. Yang. 2017. Dynamic network monitoring and control of in situ image profiles from ultraprecision machining and biomanufacturing processes. *Quality and Reliability Engineering International* 33 (8):2003–22. doi:10.1002/qre.2163.

Karrer, B., and M. E. Newman. 2011. Stochastic blockmodels and community structure in networks. *Physical Review. E, Statistical, Nonlinear, and Soft Matter Physics* 83 (1 Pt 2):016107. doi:10.1103/PhysRevE.83.016107.

Kodali, L., S. Sengupta, L. House, and W. H. Woodall. 2020. The value of summary statistics for anomaly detection in temporally evolving networks: A performance evaluation study. *Applied Stochastic Models in Business and Industry* 36 (6):980–1013. doi:10.1002/asmb.2548.

Li, D., N. Gebraeel, K. Paynabar, and A. Meliopoulos. 2021. An online approach to cyberattack detection and localization in smart grid. arXiv preprint arXiv:2102.11401.

Malko, A., C. Paris, A. Duenser, M. Kangas, D. Mollá, R. Sparks, and S. Wan. 2021. Demonstrating the reliability of self-annotated emotion data. In Proceedings of the Seventh Workshop on Computational Linguistics and Clinical Psychology: Improving Access, 45–54. doi:10.18653/v1/2021.clpsych-1.5.

McCulloh, I., and K. M. Carley. 2011. Detecting change in longitudinal social networks. Technical report. New York City, NY: Military Academy West Point NY Network Science Center (NSC).

Noble, C. C., and D. J. Cook. 2003. Graph-based anomaly detection. In Proceedings of the ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, 631–636. doi:10.1145/956750.956831.

Perry, M. B. 2019. On the detection of transitive clusters in undirected networks. *Journal of Applied Statistics* 46 (2): 364–84. doi:10.1080/02664763.2018.1491535.

Perry, M. B. 2020. An EWMA control chart for categorical processes with applications to social network monitoring. *Journal of Quality Technology* 52 (2):182–97. doi:10.1080/00224065.2019.1571343.

Perry, M. B., G. V. Michaelson, and M. A. Ballard. 2013. On the statistical detection of clusters in undirected networks. *Computational Statistics & Data Analysis* 68: 170–89. doi:10.1016/j.csda.2013.06.019.

Perry, M. B., and Z. Wang. 2020. A distribution-free joint monitoring scheme for location and scale using

individual observations. *Journal of Quality Technology* :1–18. doi:10.1080/00224065.2020.1829213.

Priebe, C. E., J. M. Conroy, D. J. Marchette, and Y. Park. 2005. Scan statistics on enron graphs. *Computational and Mathematical Organization Theory* 11 (3):229–47. doi:10.1007/s10588-005-5378-z.

Qiu, P., W. Li, and J. Li. 2020. A new process control chart for monitoring short-range serially correlated data. *Technometrics* 62 (1):71–83. doi:10.1080/00401706.2018.1562988.

Snijders, T. A., G. G. Van de Bunt, and C. E. Steglich. 2010. Introduction to stochastic actor-based models for network dynamics. *Social Networks* 32 (1):44–60. doi:10.1016/j.socnet.2009.02.004.

Tandon, A., A. Albeshri, V. Thayananthan, W. Alhalabi, and S. Fortunato. 2019. Fast consensus clustering in complex networks. *Physical Review E* 99 (4–1):042301. doi:10.1103/PhysRevE.99.042301.

Weiss, S. 2009. Privacy threat model for data portability in social network applications. *International Journal of Information Management* 29 (4):249–54. doi:10.1016/j.ijinfomgt.2009.03.007.

Wilson, J. D., N. T. Stevens, and W. H. Woodall. 2019. Modeling and detecting change in temporal networks via the degree corrected stochastic block model. *Quality and Reliability Engineering International* 35 (5):1363–78. doi:10.1002/qre.2520.

Woodall, W. H., M. J. Zhao, K. Paynabar, R. Sparks, and J. D. Wilson. 2017. An overview and perspective on social network monitoring. *IISE Transactions* 49 (3):354–65. doi:10.1080/0740817X.2016.1213468.

Xue, L., and P. Qiu. 2021. A nonparametric CUSUM chart for monitoring multivariate serially correlated processes. *Journal of Quality Technology* 53 (4):396–14. doi:10.1080/00224065.2020.1778430.

Yu, L., W. H. Woodall, and K.-L. Tsui. 2018. Detecting node propensity changes in the dynamic degree corrected stochastic block model. *Social Networks* 54:209–27. doi:10.1016/j.socnet.2018.03.004.

Yu, L., I. M. Zwetsloot, N. T. Stevens, J. D. Wilson, and K. L. Tsui. 2021. Monitoring dynamic networks: A simulation-based strategy for comparing monitoring methods and a comparative study. *Quality and Reliability Engineering*.

Zhao, M. J., A. R. Driscoll, S. Sengupta, R. D. Fricker, Jr, D. J. Spitzner, and W. H. Woodall. 2018. Performance evaluation of social network anomaly detection using a moving window–based scan method. *Quality and Reliability Engineering International* 34 (8):1699–716. doi:10.1002/qre.2364.

Zou, N., and J. Li. 2017. Modeling and change detection of dynamic network data by a network state space model. *IISE Transactions* 49 (1):45–57. doi:10.1080/0740817X.2016.1198065.