FOURIER TRAJECTORY ANALYSIS FOR IDENTIFYING SYSTEM CONGESTION

Xinyi Wu
Russell R. Barton

The Pennsylvania State University
210 Business Building
University Park, PA 16802, USA

ABSTRACT

We examine the use of the Fourier transform to discriminate dynamic behavior differences between congested and uncongested systems. Simulation continuous time statistic 'trajectories' are converted to time series for Fourier analysis. The pattern of Fourier component magnitudes across frequencies differs for congested versus uncongested systems. We use this knowledge to explore statistical process control methods to monitor nonstationary systems for transition from uncongested to congested state and vice versa. In a sense we are monitoring dynamic metamodel parameters to detect change in the dynamic behavior of the simulation. CUSUM charts on Fourier magnitudes can detect such transitions, and preliminary results suggest that in some cases detection can be more rapid than for CUSUM charts based on queue length.

1 INTRODUCTION

The design and analysis of modern discrete-event stochastic simulation has been closely tied to queuing theory; many simulation models represent a network of queues. The strength of queuing theory, and of most simulation analysis methodology, has been in deriving long-run performance measures (moments, quantiles) for stationary systems. However, many systems in real life are not stationary over time, and a good understanding of the dynamic behavior of real systems can assist managerial decision making. Simulation models can help to design that decision support: characterizing dynamic behavior in a controlled simulation setting can be used to identify effective monitoring methods for real systems. Our objective is to explore simulation analytics methods, in particular Fourier representation, for simulation trajectory data.

This study explores ways to help decision makers identify when (and which) parts of a system become congested, or move from a congested to uncongested state. This can be difficult for complex systems. Monitoring a moving average may lead to significant delays in detection. If monitoring queue length, the value must be relatively high (or low) to take into account the auto-correlated nature of the data. We examine the characteristics of simulated system trajectories, or sample paths using time series based classification methods.

Section 2 gives motivation for the proposed approach, and contrasts it with previous work in frequency domain analysis for simulation. Section 3 highlights the many technical and theoretical issues raised by this approach. Section 4 presents classification methods that we chose to test. Section 5 applies the Fourier strategy to three discrete event simulation systems. Section 6 evaluates the efficiency of different methods and suggests future explorations.

2 MOTIVATION, STRATEGY AND RELATED WORK

Our main goal is to develop a more effective approach to use dynamic data to identify or distinguish congested from uncongested queuing systems. We expect that by looking at dynamic changes of performance indicators over time, we identify whether the system is becoming congested or not. Real-time managerial decisions to reallocate system resources could then achieve less waste, and better performance.

This work is exploratory: we conduct computation experiments to explore whether Fourier representation of simulation trajectories is an effective discriminator between congested and uncongested states. There are two questions we need to answer before proceeding. First, we need a reasonable definition of system congestion. Second, we must choose dynamic performance indicators to monitor and analyze.

2.1 System Congestion

There are various definitions of system congestion in various application areas. The most common definitions for queuing systems is related to the utilization for a server. For a network of queuing systems, we consider its congestion level to be driven by the single workstation with the highest utilization.

A stable stationary system must have an arrival rate to service rate ratio strictly less than one, or a utilization strictly smaller than one. Otherwise, the queue will theoretically build up without bound. For nonstationary systems, the requirement for long-term stability is more complex, but congestion will be related to the dynamically changing utilization value. For nonstationary systems, we define system congestion as a state over a period of time rather than as a system property over all time. Choosing the length of a time period should take managerial decision making constraints into consideration.

An M/M/1 queuing system has average queue length of 81 when the system utilization is 0.9. We recognize that classifying a utilization of 0.9 as congested is context dependent, but for this study we use 0.9 utilization to assess behavior of simulated systems in a congested state.

2.2 Performance Indicator: Queue Length

For queuing systems, there are many performance measures that are affected by system congestion: the number of entities in the system or number of waiting entities, delaying time of entities, the utilization of servers, and idle/busy time of servers are a few. We use queue length at one or more servers as the focus of this study. It is feasible to get real-time values for queue length from a simulation or a real system. Delay time can only be calculated after an entity finishes service, and during which time, queue might have been significantly built up, resulting in late identification. Utilization can be estimated by calculation over time, but the estimate increases slowly after a system change. From a managerial view, the utilization is not directly observable, while the congestion level is directly observable by queue length. When utilization increases, the queue length will not increase immediately, but build up to higher peak values over time. *We hypothesize that dynamic changes in queue length might signal a transition to high utilization, perhaps sooner than simply monitoring queue length*.

Queue length is also a measurement of more realistic and practical meaning. In practice, the physical space of a queuing system is always limited, a customer may choose not to enter service if the queue length is unacceptably long.

2.3 Converting Simulation Trajectory Data into Time Series

Time series data are usually assumed to be spaced evenly in time. Discrete-event simulation trajectory data usually do not have such structure. A fundamental strategy in discrete-event simulation is to achieve computational efficiency by advancing the system clock to the next event. These events typically occur at randomly varying times.

Queue length is a continuous time statistic of a discrete event simulation. It has a piecewise constant value, which changes when an entity in the queue begins service or when an entity arrives at a busy server. An example of a queue length trajectory is shown in the upper part of Figure 1. One might assign a ‘time’ index to each change in the queue length statistic, but this would treat long sojourns at a particular queue length value the same as short sojourns at that value. Instead we propose sampling the queue length statistic at equal spacings along the trajectory. The lower part of Figure 1 shows the time series derived from the queue length trajectory data.

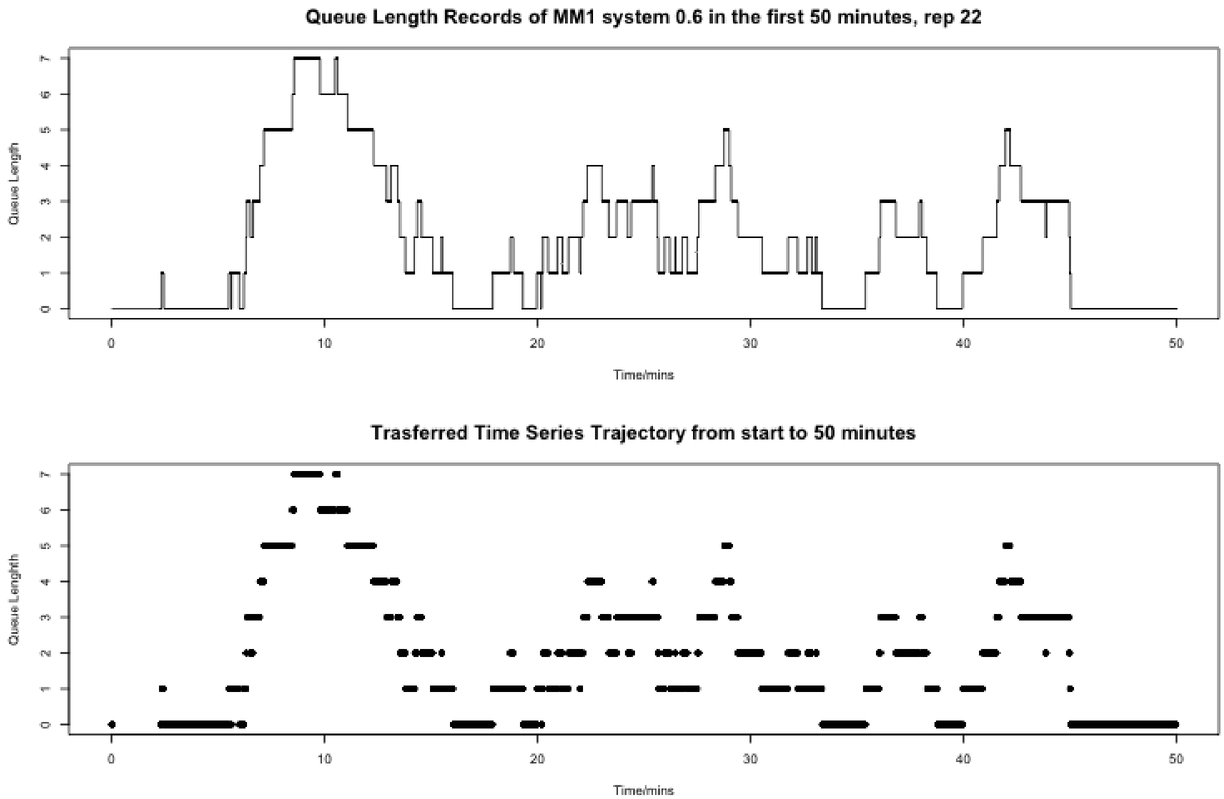

Figure 1: A sample queue length trajectory and a derived time series.

2.4 Time Series Classification Methods

The shape of the trajectory from different simulation systems can vary in characteristic ways. Some of the differences can be identified visually. Figure 2 shows queue length trajectories for M/M/1 queues with different utilizations. Congested systems have trajectories that wander, with high autocorrelation. Low-utilization trajectories exhibit a series of spikes with intervening intervals of zero queue length. While the focus here is on Fourier analysis of time series to discriminate between congested and uncongested states, there are other time series classification methods that could be employed.

2.4.1 Shapelets

Over the past several decades, there has been considerable research in time series classification. A novel method, time series shapelets, has shown potential in machine learning for image data (Ye et.al. 2010, Rakthanmanon and Keogh 2013). A shapelet is a subsequence of a time series that in some sense maximally represents a class. They are usually local patterns in a time series, characteristic highly a class of time series (Mueen, et al. 2011).

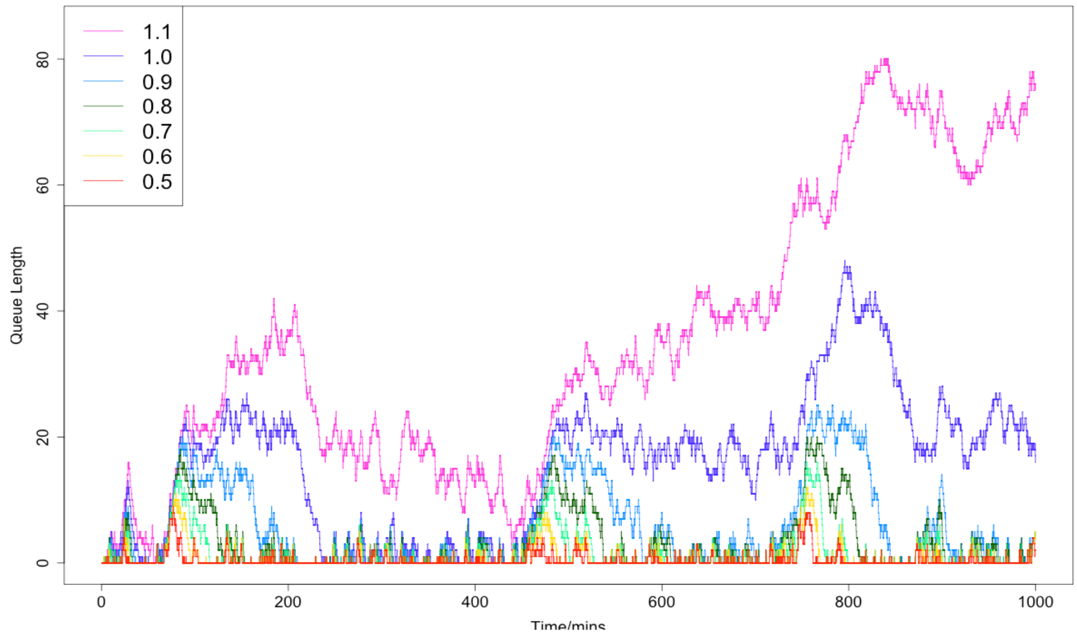

Figure 2: Trajectory of M/M/1 systems with utilizations from 0.5 to 1.1.

2.4.2 Autocorrelation Function

Congested queueing systems exhibit high autocorrelation for many performance measures, including queue length. Figure 3 shows autocorrelation plots for queue length for the multi-teller bank example in Law and Kelton (2000). The autocorrelation pattern clearly differs for the four-teller and seven-teller systems. Both have the same arrival rate, so effective utilization is much lower in the seven-teller system. We have not explored the discriminating power of functions of the autocorrelation parameters, but the autocorrelation is closely connected to the Fourier power spectrum, as we discuss below.

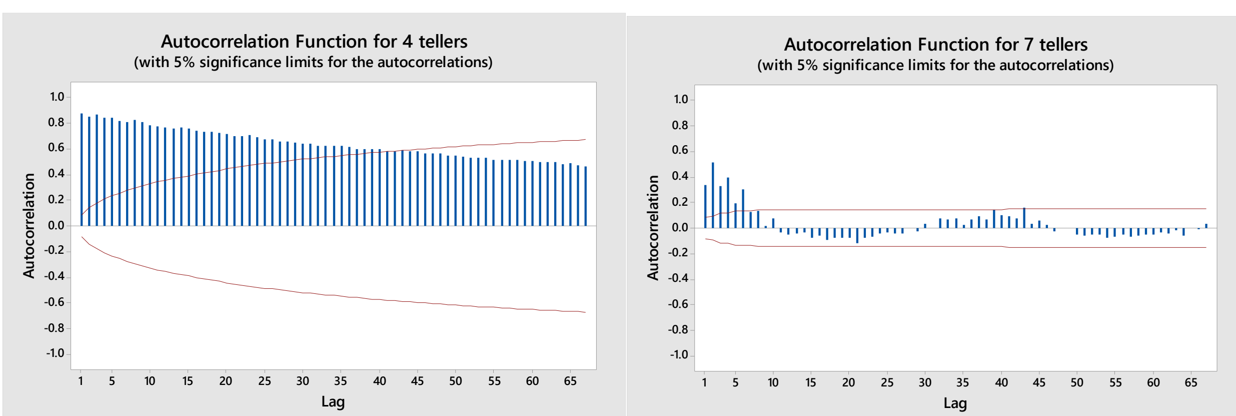

Figure 3: Autocorrelation plots for queue length for the multi-teller bank in Law and Kelton (2000).

2.4.3 Time Series Models

Time series models such as MA, ARMA and ARIMA (Box, Jenkins, and Reinsel 1994) can be fitted to time series data, and differences in fitted model coefficients could be used to discriminate between congested and uncongested systems. This discriminant can be powerful for comparing two stationary systems (as is the simple discriminator of average queue length). Difficulties arise in nonstationary situations, when changes might be not just time series model parameters, but the form of the time series model as well.

2.4.4 Fourier Decomposition

Frequency domain representation of time series have been used for many years in the economics literature. An analysis of economic time series data has found and verified a basic smooth declining shape of Fourier magnitudes exists in economic data from various sources (Granger 1966, Levy and Dezhbakhsh 2003). In spite of the presence of business cycles, seasonality and trends in economic data, the same basic shape is found regardless of these factors. Granger found that long-term fluctuation in economic variables, if decomposed into frequency components, are such that the amplitudes of the components decrease smoothly with decreasing period. In this paper, we use a spectral method, Fast Fourier Transform (FFT), to analyze simulation trajectories.

2.5 Fourier Methodology

We know that a signal can be represented in both time and frequency domains. For the same signal, its representation on each domain is linked by the Fourier transform. We will consider the discrete time Fourier transform (DTFT): although the queue length function is continuous over time, we represent it as a discrete time series as shown in Figure 1. It is possible to reconstruct a continuous-time signal in time domain by interpolating discrete signal values in the frequency domain from the output of Fourier transform. By comparing it with original trajectory, we can decide which sampling rate to use, balancing fidelity and calculation effort.

2.5.1 Representation and Reconstruction

We define an infinite signal $x(t)$ in the time domain, and $s(f)$ in the frequency domain. We can represent our signal at any instant *t* as:

$$x(t) = \sum X_n cos\{2\pi f_n t + \phi_n\} = \sum \hat{X}_n e^{j2\pi f_n t}$$

where the sum is indexed by *n*, the set of frequencies in the Fourier decomposition. The signal is composed of a series of contributions of different frequencies f_n. The size of contribution a_n, and its phase ϕ_n at *t=0* is defined by $s(f)$ at the appropriate frequency f_n. Fourier coefficients $\hat{X}_n$ could also be written in terms of magnitude and phase as

$$\hat{X}_n = A_n + jB_n$$

where the magnitude $|\hat{X}_n| = \sqrt{A_n^2 + B_n^2}/N$, the phase $\phi_n = tan^{-1}(\frac{B_n}{A_n})$, and *N* is the total number of frequencies.

In order to get the signal of a given frequency, we can see that one need two information: amplitude and phase. Actually, we can write a signal in another way as,

$$x(t) = a_0 + \sum a_n cos\{2\pi f_n t\} + \sum b_n sin\{2\pi f_n t\}$$

Because our signal is only within a finite time period $[0, T]$, the Fourier decomposition assumes the signal repeats itself outside of this interval to plus and minus infinity. Consequently it can only contain frequencies which are multiple of a fundamental frequency $f_0 = 1/T$. We can then further express our signal in the form,

$$x(t) = \sum_{n=0}^{N} a_n cos\{2\pi f_n t\} + \sum_{n=0}^{N} b_n sin\{2\pi f_n t\}$$

where a_n, b_n characterize the magnitude and phase of the *nth* frequency component corresponding to $f_n = nf_0$. Also according to the symmetry of trigonometry, we have $A_0 = a_0, 2A_n = a_n, -2B_n = b_n$.

Figure 4 gives a graphical representation of our overall methodology: a simulation trajectory is transformed into a time series with time intervals of size δ, and magnitudes of the Fourier components are used to discriminate the dynamic behavior of uncongested vs. congested systems.

Since the output of DTFT returns complex numbers representing Fourier coefficients $\hat{X}_n$ at each frequency, we can calculate the corresponding value of a_n, b_n and reproduce the signal value by interpolating them into above equation at discrete time points. In fact, reconstruction using a subset of the full set of Fourier components will approximate the observed time series. *Thus Fourier representations can be thought of as a metamodel for the dynamic behavior of a discrete-event simulation*.

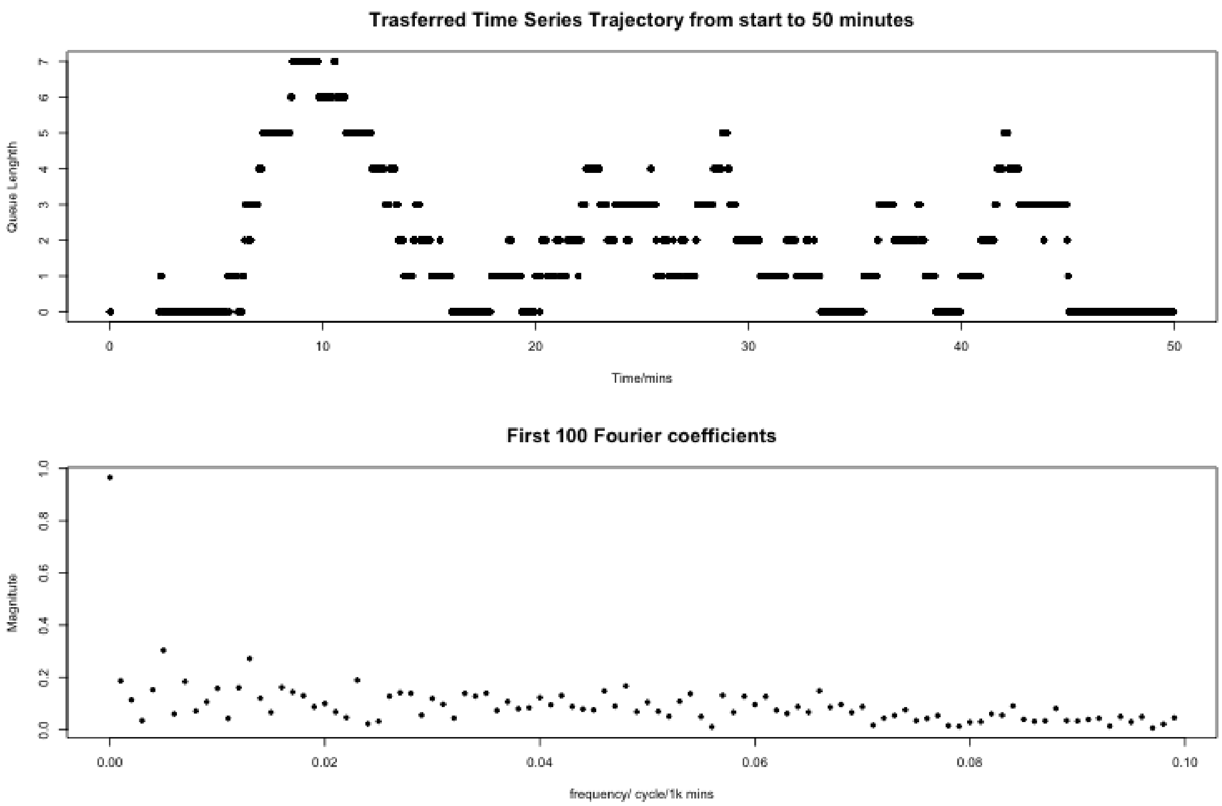

Figure 4: Moving from time series to Fourier magnitudes for the M/M/1 queue trajectory in Figure 1.

2.5.2 Discrimination by Magnitude Components at Various Frequencies

For our analysis, we compare the magnitude values at different frequencies for different systems. From trajectory plots of M/M/1 systems with different utilization levels. From Figure 2 above, We hypothesize that uncongested systems have more small spikes connected by zero value line segments and congested systems have larger positive oscillations over longer time periods. So in the output of our DTFT, we expect to observe magnitude differences at both low and high frequencies.

2.6 Distinction from Prior Frequency Domain Simulation Research

The Fourier approach that we take is distinct from the frequency domain sensitivity analysis by Schruben and Cogliano (1981) and the many subsequent papers. In that research, simulation input parameters were deliberately varied periodically to determine sensitivity of simulation output to the input parameters. By varying different input parameters at (carefully chosen) different frequencies, sensitivities for multiple parameters could be determined from only two simulation runs. We do not seek sensitivity information, and do not deliberately vary parameters for that purpose, rather we seek to characterize a system's dynamic behavior by its Fourier signature, for use in discriminant analysis and process monitoring.

Further, the index used for the Fourier calculations in prior frequency domain work was typically an entity index. Input parameters were varied with entity index, and discrete-time output statistics were analyzed based on the index of the associated entities. Jacobson, Morrice and Schruben (1988) did examine using simulation clock tie for the driving frequencies, but still analyzed output using an entity index.

Hazra and Park (1994) used fixed-time increments by binning arrival and departure events into time intervals. They used this approach to characterize the frequency response of an M/G/1 queue. Though this was again applied in the context of deliberate periodic variation of input parameters, their findings included a result that is useful in our setting: an M/G/1 queue acts as a low-pass filter for temporal variations in arrivals, and the pass-band decreases with increases in system utilization. *In certain queueing systems, then, we might expect the Fourier representation to have relatively less high-frequency content if the system is congested*.

Although not in a frequency domain context, other dynamic characterization of queues might also inform schemes for detecting changes in congestion. See for example Eick, Massey, and Whitt (1993).

3 TECHNICAL ISSUES

Our exploratory investigation considered two issues: i) can patterns in the frequency domain discriminate the dynamic behavior of congested vs. uncongested systems, and ii) can statistical process control methods based on these findings be effective. In order to conduct computation experiments, several technical issues had to be examined.

3.1 Potential Problems of the Time Interval Selection

In order to apply DTFT to an output sample path of queue length changes over time, it must be transformed into time series data. Thus, we need to choose a time sampling interval small enough to enable capture of key dynamic features. If we select the minimal positive time interval observed between events in the original data as the time interval, the total number of points in the time series would be computationally intractable for some systems. For example, for the multi-teller bank (MTB) system with 5 tellers, 0.6 utilization and fixed single server mean service rate at 1 per minute, for a 480-minute sample trajectory, the minimal positive time interval in a particular replication was 0.000075 minute, resulting in more than 64 million time series points in 480 minutes, while the actual number of queue length changes in that sample path was only 803. And most of the calculation in our transform seems unnecessary for uncongested systems, with queue length tending to remain at 0 for long periods. On the other hand, if we use a larger interval, we will fail to capture some queue length changes that occurred in the trajectory data.

3.1.1 Time Interval Selection -- 10^{-k} percentile value of inter-arrival/service time distribution

For the simulation systems in this paper, both inter-arrival and service times have been assumed to be exponentially distributed. We considered using a 10^{-k} percentile time of corresponding exponential distribution as interval time to generate time series data.

There are two main reasons for this. First, since the queue length in our system will only be possible to change when entity arrives, departs/moves between servers/workstations, the time interval between changes will follow exponential distributions. And by selecting the 10^{-k} percentile value of the interarrival time distribution, the probability that the next arrival is in a different interval is $1\text{-}10^{-k}$. Second, this 10^{-k} percentile value can be easily calculated and bounds the number of intervals likely to have more than one event.

In this paper, we will select 10^{-3} percentile value. This means, we expect to have around 0.001 possibility to include more than one queue length change within new time interval, and around $(0.001)^2$ possibility to include consecutive two changes. If an inclusion is detected, we shift the event to the next interval and shift all later trajectory transition points by the same amount. This allows keeping all queue change events in the time series representation, with only a small perturbation to the original data values.

In order to speed up the DTFT calculation via the FFT algorithm (Cooley and Tukey 1965) and to enable us to compare outputs of DFFT under different system settings, we further adjust the interval value, so that in total number of generated points is a power of 2.

3.1.2 Time Interval Selection – Fidelity

When using the adjusted distribution 0.001 percentile value as time interval resulted in total points at 2^{20}, or 1,048,576 to the MTB. Is this enough? There are two additional justifications that we considered for the time interval choice. First, the importance of short-duration queue length changes from a managerial perspective. For the bank teller example, queue length changes of less than a second in duration probably have little meaning to the manager, and we suspect are not critical in determining system congestion.

Second, a sufficient sampling rate guarantees the fidelity of our DTFT. Combined with minimal managerially interesting time intervals, this can lead to a minimal time interval choice. According to Nyquist-Shannon sampling theorem, the necessary sampling frequency is two times of the maximum signal frequency of our simulation trajectory. If the minimal interesting interval is one second, a 480-minute, 28,800-second simulation would require sampling at least 57,600 times, a value between 2^{15} and 2^{16}.

3.1.3 Signal Reconstruction as a Check

Due to the nature of queuing system, the trajectory of queue length changes consists of various square waves. Thus we have discontinuity in our trajectory, and the FFT output from a relatively low sampling rate would actually represent a signal with strong ringing artifacts – the Gibbs phenomenon. In order to better separate systems with different congestion levels, we need to guarantee a more precise transform outcomes by sampling at higher than the Nyquist frequency for sine waves with period equal to the minimal interesting time interval.

To check that a proposed sampling rate produces adequate fidelity, one can examine the fidelity of the reconstructed signal (as described in section 2.5.1). We will illustrate this in section 3.2.

3.2 Frequency Component Selection for Discrimination

In order to better differentiate systems, we need to select the Fourier coefficient magnitude value(s) which could best distinguish system trajectory characteristics. Define the magnitude value at frequency f_n as C_n. We proposed to look into three combinations: The first magnitude value C_1(this is identically the mean queue length), the average value of the first k_1 but the first magnitude $\frac{\sum_{i\in[2,k_1]} C_i}{k_1-1}$ and average magnitude value over selected frequency range $\frac{\sum_{i\in[k_2,k_3]} C_i}{k_3-k_2}$.

Intuitively, the trajectory of uncongested system would normally have more small spikes, with small width and height, while congested system tend to have wandering and higher levels, corresponding to lower frequencies. In general, a congested system also has higher magnitude values over the whole spectrum. So we will first use the average of low frequency components as one of our indicators. And the selection of k_2 and k_3 will be based on examination of spike interval values existed in uncongested system trajectory. We will find the interval value, where more than half of the spike interval values falls between it and zero. And k_2 and k_3 equals to trajectory length T divided by the length range of small peaks in the trajectory, which represents the possible frequency range where their magnitude value may appear. The fundamental frequency $f_0 = 1/T$ stands for how frequently one component will appear within time T.

3.3 Statistical Process Control for Autocorrelated Data (SPC)

Traditionally, control charts are used to continuously monitor system performances over time with specified upper and lower limits. The existence of a large amount of historical data set and the assumption that process observations are i.i.d are typical conditions to effectively and efficiently implement control charts (Snoussi et al. 2005). For this paper, the systems under study are often networks of queues. Output will exhibit significant autocorrelation.

One strategy for monitoring autocorrelated time series is to fit a time series model such as those described above during a calibration period and then monitor the residuals via CUSUM or EWMA (Apley and Shi 1999; Runger, Willemain, and Prabhu 1995; Nenes and Taragas 2005). Our concern with this approach is the uncertainty in model form. Apley (2002) discusses handling uncertainty in parameter values, but potential epistemic uncertainty remains. This is of special concern if the nonstationarity evidenced in unknown changes in model form.

The cumulative sum (CUSUM) chart was specifically proposed and tested for monitoring queue performance metrics of an M/M/1 queuing system. Chen and Zhou (2015) demonstrated that CUSUM was effective for monitoring estimated system parameters inter-arrival rate and service rate.

4 EXPERIMENTS

To test our proposed methods, we examined three different queuing systems with different levels of complexity. All three simulations used two different combinations of key parameter settings in order to achieve target levels of utilization, each setting with 50 replications. The systems were an M/M/1 queue with utilizations of 0.6 and 0.9, and the multi-teller bank and job shop models from Law and Kelton (2000). The multi-teller bank simulation used five tellers and utilizations of 0.6 and 0.9. For each system, two sets of experiments were performed: first, Fourier analyses of entire trajectories were computed, and coefficient magnitudes were examined for their ability to discriminate between congested and uncongested systems. Generally we found that Fourier magnitudes could be used to distinguish between congested and uncongested systems, although this was weakest for the M/M/1 queue. Figure 5 shows magnitude comparison for congested and uncongested trajectories for the multi-teller bank simulation. There is little overlap in the average magnitude of Ck_2-Ck_3 coefficients.

4.1 Preliminary SPC Results

For examining the effectiveness of Fourier methods for detecting change in congestion we employed the common testing regime of a fixed change in congestion (utilization) following the calibration period. Figures 6-9 show boxplots of the time until signal for each of the SPC methods that we examined for each of the three simulation models. Each boxplot is based on 50 replications. We see that the dynamic behavior of queue length, as characterized by Fourier coefficient, can be used to detect changes in system congestion.

These preliminary results suggest that Fourier-based control charting may result in improvements over simple monitoring of queue length, when systems move from congested state to uncongested state.

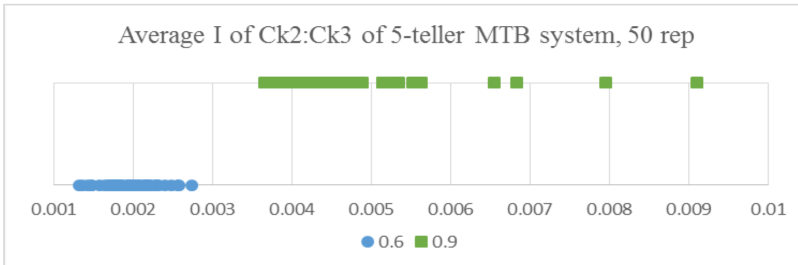

Figure 5: Fourier coefficient magnitudes for congested/uncongested MTB systems.

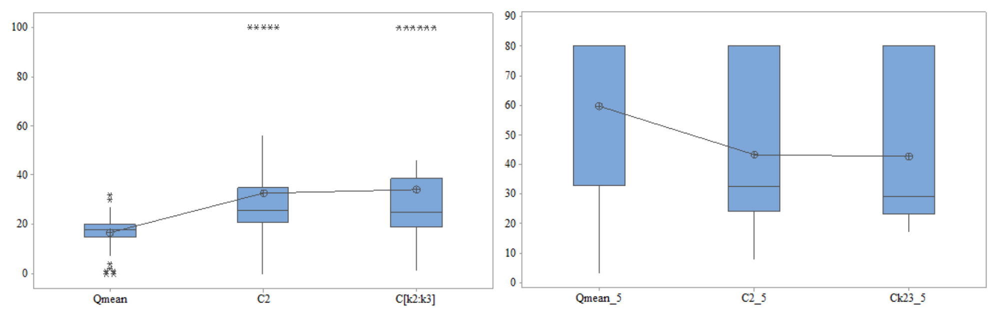

Figure 6: SPC performance for the M/M/1 queue. Asterisks are replications without detecting in the experiment window. Vertical axis is time to out-of-control signal. Left: increase in congestion. Right: decrease.

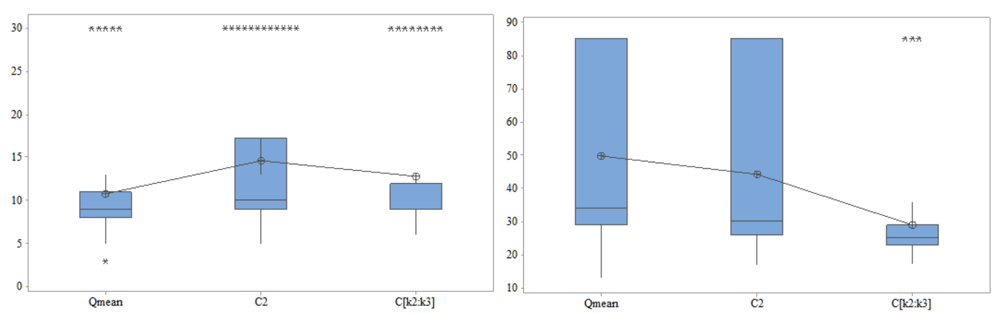

Figure 7: SPC performance to MTB (increase in congestion) and Job Shop (decrease).

5 SUMMARY

We have two interesting findings in this investigation. First, it appears that frequency domain methods can be used to distinguish the dynamic behavior of congested versus uncongested systems. Second, this power can be applied in a monitoring situation, to allow decision makers to react to changes in system dynamics. Simulation provides a mechanism to test such dynamic detection (and control) procedures before they are implemented on a real system. But these findings are preliminary; much remains to be done.

First, we have not attempted to optimally select i) sampling frequency, ii) identification of Fourier components with maximal discriminatory power. Second, discriminatory power may be magnified by retaining the multivariate nature of the Fourier coefficients. In particular, using a multi-variate SPC method, may result in better and more accurate SPC signaling. And for all of these opportunities, it will be important to provide theoretical and empirical performance estimates, to be sure that the methods merit adoption.

Finally, for the SPC method implemented here, we move the time window of Δt width by Δt to construct the next observation data for control charting. This corresponds to nonoverlapping batches. It may be possible to improve responsiveness using overlapping samples for construction of Fourier magnitudes, for SPC.

REFERENCES

Apley, D. W. 2002. "Time Series Control Charts in the Presence of Model Uncertainty." *Transactions of the ASME* 124: 891–898.

Apley, D. W., and J. Shi. 1999. "The GLRT for Statistical Process Control of Autocorrelated Processes." *IIE Transactions* 31: 1123–1134.

Box, G. B. P., G. Jenkins, and G. Reinsel. 1994. *Time Series Analysis, Forecasting, and Control*, 3rd ed. Englewood Cliffs, NJ: Prentice-Hall.

Chen, N., and S. Zhou. 2015. "CUSUM Statistical Monitoring of M/M/1 Queues and Extensions." *Technometrics* 57: 245–256.

Cooley, J. W. and J. W. Tukey. 1965. "An Algorithm for the Machine Calculation of Complex Fourier Series". *Mathematics of Computation* 19: 297–301.

Eick, S. G., W. A. Massey, and W. Whitt. 1993. "The Physics of the Mt/G/∞ Queue." *Operations Research* 41: 731–742.

Granger, C. W. J. 1966. "The Typical Spectral Shape of an Economic Variable." *Econometrica* 34: 150.

Hazra, M. M. and S. K. Park. 1994. "Characterizing a Nonstationary M/G/1 Queue using Bode Plots." In *Proceedings of the 1994 Winter Simulation Conference*, edited by J. D. Tew, S. Manivannan, D. A. Sadowski, and A. F. Seila, 377-382. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Jacobson, S. H., D. Morrice, and L. W. Schruben. 1988. "The Global Simulation Clock as the Frequency Domain Experiment index." In *Proceedings of the 1988 Winter Simulation Conference*, edited by M. Abrams, P. Haigh and J. Comfort, 558-563. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Keogh, E., and S. Kasetty. 2003. "On the Need for Time Series Data Mining Benchmarks: A Survey and Empirical Demonstration." *Data Mining and Knowledge Discovery* 7: 349–371.

Law, A. M., and W. D. Kelton. 2000. *Simulation Modeling & Analysis*. 3rd ed. New York: McGraw-Hill, Inc.

Levy, D., and H. Dezhbakhsh. 2003. "On the Typical Spectral Shape of an Economic Variable." *Applied Economics Letters* 10: 417.

Mueen, A., E. Keogh, and N. Young. 2011. "Logical-Shapelets: An Expressive Primitive for Time Series Classification." In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1154–1162. KDD '11. New York, NY, USA: ACM.

Nenes, G. and Tagaras, G., 2005. The CUSUM chart for monitoring short production runs. In *Proceedings of 5th International Conference on Analysis of Manufacturing Systems–Production Management*. Zakynthos Island, Greece, 43-50.

Rakthanmanon, T., and E. Keogh. 2013. "Fast Shapelets: A Scalable Algorithm for Discovering Time Series Shapelets." In *Proceedings of the 2013 SIAM International Conference on Data Mining*, 668–676. Society for Industrial and Applied Mathematics.

Runger, G. C., T. R. Willemain, and S. Prabhu. 1995. "Average Run Lengths for Cusum Control Charts Applied to Residuals." *Communications in Statistics: Theory & Methodology* 24: 273–282.

Schruben, L. W. and V. J. Cogliano. 1981. "Simulation Sensitivity Analysis: A Frequency Domain Approach." In *Proceedings of the 1981 Winter Simulation Conference*, edited by T. I. Oren, C. M. Delfosse and C. M. Shub, 455-459. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Snoussi, A., M. El Ghourabi, and M. Limam. 2005. "On SPC for Short Run Autocorrelated Data." *Communications in Statistics: Simulation & Computation* 34: 219–234.

Ye, L., and E. Keogh. 2010. "Time Series Shapelets: A Novel Technique That Allows Accurate, Interpretable and Fast Classification." *Data Mining and Knowledge Discovery* 22 (1-2): 149–82.

AUTHOR BIOGRAPHIES

XINYI WU is a graduate student in the Department of Industrial and Manufacturing Engineering at Penn State University. She will finish her Master's degree in Industrial Engineering and Operations Research in the May of 2016. She is interested in simulation, supply chain modeling and optimization. Her email address is wuxinyi8@gmail.com.

RUSSELL R BARTON is professor of supply chain and information systems and professor of industrial engineering at the Pennsylvania State University. He currently serves as senior associate dean for research and faculty in the Smeal College of Business. He received a B.S. degree in electrical engineering from Princeton University and M.S. and Ph.D. degrees in operations research from Cornell University. He serves as the I-Sim representative on the INFORMS Subdivisions Council and chairs the QSR Advisory Board. He is a senior member of IIE and IEEE and a Certified Analytics Professional®. His research interests include applications of statistical and simulation methods to system design and to product design, manufacturing and delivery. His email address is rbarton@psu.edu.