Proceedings of the 2006 Winter Simulation Conference
L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, eds.

MONITORING VARIABILITY OF AUTOCORRELATED PROCESSES USING STANDARDIZED TIME SERIES VARIANCE ESTIMATORS

Seong-Hee Kim

H. Milton Stewart School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0205, U.S.A.

ABSTRACT

We consider the problem of monitoring variability of autocorrelated processes. This paper combines variance estimation techniques from the simulation literature with a statistical process control chart from statistical process control (SPC) literature. The proposed SPC method does not require any assumptions on the distribution of the underlying process and uses a variance estimate from each batch as a basic observation. The control limits of the chart are determined analytically. The proposed chart is tested using stationary processes with both normal and non-normal marginals.

1 INTRODUCTION

Statistical process control (SPC) charts are widely used to detect shifts in the parameters of monitoring processes. Recently, the problem of monitoring an autocorrelated process has been received a lot of attention from the SPC community and majority of effort has been focused on developing distribution-based methods. Distribution-based SPC charts require the in-control underlying process to follow a specific probability distribution model, or certain characteristics of the process, such as the autocorrelation structure, to be known. For distribution-based SPC charts that monitor the mean of autocorrelated processes, see Kim et al. (2006a) for a detailed review. Distribution-based SPC charts are often criticized from the fact (a) that the underlying assumptions may be violated, resulting that the charts may not work as advertised and (b) that their control limits are often determined by trial-and-error which is sometimes inconvenient in practical applications due to time or availability of data. These limitations could be overcome by distribution-free SPC charts. There are a few distribution-free methods for monitoring mean, which include Johnson and Bagshaw (1974), Runger and Willemain (1995), and Kim et al. (2006a, b).

While the problem of monitoring variability of an autocorrelated process is as important as that of monitoring mean, little work has been done in the former problem compared to the latter problem. Reynolds and Lu (1997) and Lu and Reynolds (1999) present SPC charts for monitoring variance of autocorrelated processes under the assumption that the underlying process follows an ARMA(1,1) process. Cook and Chui (1998) introduce neural network techniques for monitoring parameters of autocorrelated processes. Chiu, Chen, and Lee (2001) and Cook, Zobel, and Nottingham (2001) apply this method to monitor mean and variance of an autocorrelated process, assuming that the underlying process follows an AR(1) model. Their SPC methods with the neural network techniques may be extended to distribution-free methods. However, such work has not been proposed yet and the methods require training on both in-control and *out-of-control* data unlike usual SPC charts that require training on the in-control process only. To our best knowledge, there is no explicitly proposed distribution-free SPC chart for the purpose of detecting shifts in variability of autocorrelated processes in SPC literature.

In this paper, we develop distribution-free SPC charts for monitoring variability of autocorrelated processes by combining techniques from simulation and SPC literature. A good measure for variability of autocorrelated processes is so called the asymptotic variance parameter which is basically the sum of covariances at all lags. There are a number of distribution-free techniques for estimating the asymptotic variance parameter in simulation literature; see Alexopoulos, Goldsman, and Serfozo (2005). Among those variance estimation techniques, we take a variance estimator that can be computed from one batch rather than several batches, and those estimates will be used as basic observations for SPC charts. The expectation of the variance estimates from the in-control process is approximately equal to the in-control asymptotic variance parameter if a batch size is large enough. Therefore, the problem of monitoring variability becomes monitoring the mean of variance estimates, and we can apply existing distribution-free SPC charts originally developed for monitoring mean.

In Section 2, we give notation and assumption for the paper and review a variance estimation technique for the asymptotic variance parameter. A new distribution-free SPC method for monitoring variability is presented in Section 3. Experimental studies based on AR(1) processes and M/M/1 queueing models are given in Section 4, followed by conclusions in Section 5.

2 BACKGROUND

In this section, we define notation and assumptions on output data from the in-control monitoring process. We also give a review of a distribution-free estimator for the asymptotic variance parameter.

2.1 Notation and Assumptions

Let $\{Y_i : i = 1,2,\ldots\}$ denote the discrete-time monitoring process that has a steady-state distribution with marginal mean $\mathrm{E}[Y_i] = \mu$ and marginal variance $\mathrm{Var}[Y_i] = \sigma^2$. Also, let $\bar{Y}(n)$ denote the sample mean of the first n observations. Then, the standardized CUSUM, $\mathcal{C}(t)$ is defined as

$$\mathcal{C}(t) \equiv \frac{\sum_{j=1}^{\lfloor nt \rfloor} Y_j - nt\mu}{\Omega\sqrt{n}}, \quad 0 \le t \le 1, \tag{1}$$

where $\lfloor\cdot\rfloor$ is the "floor" (greatest integer) function and Ω^2 is the asymptotic variance constant, defined as

$$\Omega^2 \equiv \lim_{n\to\infty} n\,\mathrm{Var}(\bar{Y}(n)) = \sum_{\ell=-\infty}^{\infty} \mathrm{Cov}\big[Y_i, Y_{i+\ell}\big],$$

where we assume that $0 < \Omega^2 < \infty$. The in-control asymptotic variance parameter is denoted by Ω_0^2.

Let $\mathcal{W}(\cdot)$ denote a standard Brownian motion process on $[0,\infty)$ so that $\mathcal{W}(t)$ is normally distributed with $\mathrm{E}[\mathcal{W}(t)] = 0$ and $\mathrm{Cov}[\mathcal{W}(s),\mathcal{W}(t)] = \min\{s,t\}$ for $s,t \in [0,\infty)$. The random function $\mathcal{C}(t)$ is an element of the Skorohod space $D[0,1]$, i.e., the space of functions on $[0,1]$ that are right-continuous and have left-hand limits (Chapter 3 of Billingsley 1968). We restrict our interest on processes that satisfy the following assumption which is called a Functional Central Limit Theorem (FCLT) (see Billingsley 1968, Chapter 4).

Assumption 1 *(FCLT) There exist finite real constants μ and $\Omega^2 > 0$ such that the probability distribution of $\mathcal{C}(t)$ over $D[0,1]$ converges to that of $\mathcal{W}(\cdot)$ for $t \in [0,1]$, as $n \to \infty$. Formally,*

$$\mathcal{C}(t) \underset{n\to\infty}{\overset{D}{\longrightarrow}} \mathcal{W}(t), \quad 0 \le t \le 1,$$

where $\underset{n\to\infty}{\overset{D}{\longrightarrow}}$ denotes convergence in distribution as $n \to \infty$.

Further, we assume that for every $t \in [0,1]$, the family of random variables $\{\mathcal{C}^2(t) : n = 1,2,\ldots\}$ is uniformly integrable (see Billingsley 1968, Chapter 5).

For autocorrelated processes, a good measure for variability is the asymptotic variance parameter Ω^2 and this paper focuses on monitoring changes in Ω^2.

2.2 Variance Estimators

In this subsection, we review one estimator for the variance parameter Ω^2. The simplest and most popular estimator for Ω^2 is probably the nonoverlapping batch means estimator that first forms batch means (i.e., sample averages of contiguous, but autocorrelated, observations) and computes the usual sample variance of the batch means, provided that the batch size is large enough to ensure that batch means are approximately independent and identically distributed (i.i.d.). For the nonoverlapping batch means method, one needs at least two batches to get one variance estimate and this is not desirable for the purpose of monitoring. Runger and Willemain (1995) propose a method of using batch means as basic observations for monitoring mean of autocorrelated processes. The method is distribution-free, but it was pointed out that the method may delay legitimate out-of-control alarms for highly correlated processes or large shifts. This problem will be more serious if a basic observation of an SPC chart requires at least two batches. For this reason, we do not consider the nonoverlapping batch means method. Instead, we consider an estimator whose estimate can be computed from one batch. Any estimator based on standardized time series (see Schruben 1983 for the definition of standardized time series) can generate an estimate from one batch. Alexopoulos, Goldsman, and Serfozo (2005) provide a review of various standardized time series variance estimators. In this paper, we consider the CvM estimator only. Kim (2006) consider other variance estimators for monitoring variability of autocorrelated processes.

The method we discuss below relies on Assumption 1 (together with mild moment and mixing conditions) and produces an asymptotically consistent estimator of the asymptotic variance parameter, Ω^2.

We define the area under the square of the STS and its limiting functional as

$$C(g;n) \equiv \frac{1}{n}\sum_{k=1}^{n} g(\frac{k}{n})\Omega^2 T_n^2(\frac{k}{n})$$

and

$$C(g) \equiv \int_0^1 g(t)\Omega^2\mathcal{B}^2(t)\,dt,$$

respectively, where $g(t)$ is a weighting function normalized so that $\mathrm{E}[C(g)] = \Omega^2$ and $g''(t)$ is continuous and bounded on $[0,1]$. Under mild assumptions, the CMT implies that $C(g;n) \xrightarrow[n\to\infty]{\mathrm{D}} C(g)$, and we call $C(g;n)$ the *weighted Cramér–von Mises* (CvM) estimator for Ω^2.

Goldsman, Kang, and Seila (1999) show that under Assumption 1 and some mild moment and mixing conditions, the weighted CvM estimator can be first-order unbiased. For example, the choice of $g(t) \equiv -24+150t-150t^2$ results in $\mathrm{E}[C(g;n)] = \Omega^2 + o(1/n)$. The variance of the weighted CvM estimator depends on the weighting function $g(t)$. With the choice of $g(t) = -24+150t-150t^2$, we get

$$\mathrm{Var}(C(g)) = 1.729\Omega^4. \tag{2}$$

There is another weighting function for the CvM estimator that results in a first-order unbiased estimator with a slightly smaller variance. However, Goldsman, Kang, and Seila (1999) show that the CvM estimator with $g(t) = -24+150t-150t^2$ is more reliable, and we employ this weighting function for the CvM estimator in this paper.

For the purpose of monitoring variability of autocorrelated processes, we get an estimate from each batch and use those estimates as basic observations for an SPC chart. Throughout the paper, V_i represents a CvM variance estimate from the ith batch with the weighting functions $g(t) = -24+150t-150t^2$.

3 MONITORING VARIABILITY

In this section, we discuss how to monitor changes in variability of autocorrelated processes. As discussed in the previous section, variance estimation techniques based on standardized time series is distribution-free methods for estimating the asymptotic variance parameter Ω^2 based on batching. If we take variance estimates for Ω^2 from batches as basic observations, the problem of monitoring the variability of $Y_1,Y_2,\ldots,$ becomes that of monitoring the mean of $V_1,V_2,\ldots$. There are a number of distribution-free SPC charts for monitoring mean. By combining a distribution-free variance estimation technique with a distribution-free SPC chart for monitoring mean, one can come up with a distribution-free SPC chart for monitoring variability.

There are distribution-free SPC charts that do not assume any distribution on basic observations, and they include Johnson and Bagshaw (1974) and Kim et al. (2006a, b). Among those procedures, the procedure due to Kim et al. (2006b) shows the most efficient performance. Therefore we consider their distribution-free Tabular CUSUM chart in this paper and propose the distribution-free Tabular CUSUM for Variability (DFTCV) chart (see Kim 2006 for other SPC charts). The procedure requires the variance of basic observations which is the variance of $V_1,V_2,\ldots$ in our case. We denote the variance of CvM estimates by Ψ^2 and its in-control value by Ψ_0^2.

In reality the in-control Ω_0^2 and Ψ_0^2 should be estimated from a training data set. Estimating Ψ_0^2 would require batching of already batched data if the estimates from batches are autocorrelated, which may require a huge data set. On the other hand, if a FCLT approximately holds and variance estimates are independent, then we can estimate Ω_0^2 from a training data set and use (2) to get Ψ_0^2. For this reason, we choose a batch size m that makes V_i approximately independent and ensures an approximate FCLT. Then the DFTCV chart is presented below:

1. Set $K=0.1$, $\Psi_0 = 0.1\sqrt{1.729}\Omega_0^2$, and a target two-sided ARL_0 in terms of raw observations. Then, calculate H, the solution to the equation:

$$\frac{\Psi_0^2}{2K^2}\left\{\exp\left[\frac{2K(H+1.166\Psi_0)}{\Psi_0^2}\right]-1 - \frac{2K(H+1.166\Psi_0)}{\Psi_0^2}\right\} = \frac{2\mathrm{ARL}_0}{m}.$$

2. Raise an out-of-control alarm after the ith batch if $S^+(i) \geq H$ or $S^-(i) \geq H$ where $S^+(0)=0, S^-(0)=0$; and for $i=1,2,\ldots$, $S^+(i) = \max\{0,\ S^+(i-1)+(V_i-\Omega_0^2)-K\}$ and $S^-(i) = \max\{0,\ S^-(i-1)-(V_i-\Omega_0^2)-K\}$.

One critical issue on implementing the DFTCV chart is to determine a batch size. Basically a batch size m should be large enough to ensure (i) that variance estimates V_i are approximately independent and (ii) that a FCLT approximately holds. For the first condition, we check if batch means are approximately independent. For the second condition, we check if a batch size is approximately normally distributed. Lada and Wilson (2005) employ the von Neumann test (von Neumann 1941) for the independence test of batch means and the Shapiro-Wilk test (Shapiro and Wilk 1965) for the normality test. We modify the algorithm of Lada and Wilson (2005) to determine an appropriate batch size m and the algorithm is presented in Kim (2006).

4 EXPERIMENTS

We test the performance of the proposed charts on stationary processes with both normal and non-normal marginals. For normal marginals, we employ a stationary first-order autoregressive (AR(1)) process. For non-normal marginals, the queue waiting times observed in an $M/M/1$ queue are considered.

4.1 AR(1) Processes

An AR(1) process is defined as follows:

$$Y_j = \mu + \varphi(Y_{j-1} - \mu) + \varepsilon_j \quad \text{for} \quad j = 1, 2, \ldots, \qquad (3)$$

where: (i) $\{\varepsilon_j : j = 1, 2, \ldots\} \overset{\text{i.i.d.}}{\sim} \mathrm{N}(0, \sigma_\varepsilon^2)$; (ii) we take $-1 < \varphi < 1$ to ensure that (3) defines a stationary AR(1) process; and (iii) we take $Y_0 \sim \mathrm{N}(\mu, \sigma^2)$ to ensure that the process $\{Y_j\}$ starts in steady-state operation.

For the AR(1) process (3), the marginal variance is

$$\sigma^2 = \frac{\sigma_\varepsilon^2}{1-\varphi^2};$$

the lag-ℓ covariance is

$$\mathrm{Cov}(Y_i, Y_{i+\ell}) = \sigma^2 \varphi^{|\ell|} = \frac{\sigma_\varepsilon^2 \varphi^{|\ell|}}{1-\varphi^2}$$

for $\ell = 0, \pm 1, \pm 2, \ldots$, and the variance parameter is

$$\Omega^2 = \sigma^2 \left(\frac{1+\varphi}{1-\varphi}\right) = \frac{\sigma_\varepsilon^2}{(1-\varphi)^2}.$$

The parameters of the in-control AR(1) process are denoted by σ_0^2, $\sigma_{0\varepsilon}^2$, Ω_0^2, and φ_0. Specifically, σ_0^2 is set to one for all in-control process; therefore, $\sigma_{0\varepsilon}^2 = 1 - \varphi_0^2$. Let Ω^2 denote the asymptotic variance parameter of an in-control or out-of-control process. Then the ratio of Ω^2/Ω_0^2 varies over 1, 1.5, 2, 2.5, 3, 4, and 10. The coefficient φ_0 is set to 0.25, 0.3, 0.5, 0.7, 0.9, and 0.95.

There are two sources for a shift in Ω^2: (i) changes are caused by a shift in either σ^2 or σ_ε^2 and (ii) changes are caused by a shift in φ. When the ratio $\Omega^2/\Omega_0^2 = c$ and the change is caused by a shift in σ^2 or σ_ε^2, the shifted σ^2 is set to c, and therefore the shifted $\sigma_\varepsilon^2 = c(1-\varphi_0^2)$. Similarly, when the ratio $\Omega^2/\Omega_0^2 = c$ and the change is caused by a shift in φ, the shifted φ is set to

$$\frac{(1+c)\varphi_0 + (c-1)}{(1+c) + (c-1)\varphi_0}.$$

The target ARL_0 is set to 10,000. Notice that 10,000 is considerably large compared to 370 which has been a norm for SPC charts for i.i.d. data. However, autocorrelation into the data is mainly caused by the introduction of automated testing devices for monitoring with high sampling frequency. This frequent automated sampling is becoming more and more common. Therefore, an ARL of 370 might correspond to only 370 minutes (or even seconds) of process operation, which is woefully inadequate in many applications. We want a much longer time between false alarms, and in-control ARLs in excess of 10,000 are considered to be desirable in a number of papers that deal with autocorrelated processes (for example, see Yashchin 1993 and Runger and Willemain 1995).

Table 1 shows the experimental results when the shift in variability is caused by a shift in σ^2 or σ_ε^2. The batch size m is determined by an algorithm presented in Kim (2006). We found that a batch size determined by the algorithm works well for the DFTCV chart in a sense that it results in actual ARL_0 close to the target value 10,000.

Table 2 shows the experimental results when the shift is caused by a shift in φ. Notice that for the DFTCV chart ARL_1 decreases as Ω^2/Ω_0^2 increases, and then *increases* when $\Omega^2/\Omega_0^2 = 10$. When φ is shifted to a large value, it affects the autocorrelation structure of the monitoring process, and observations become more strongly correlated. Therefore, the batch size m that used to work well for the in-control process is not large enough any more for the out-of-control process, allowing high bias to sneak into variance estimates and resulting low-biased estimates. We conjecture that this is why ARL_1 for $\Omega^2/\Omega_0^2 = 10$ is larger than that for a smaller shift.

4.2 M/M/1 Queue Waiting Times

For the stationary process with non-normal marginal, we consider the queue waiting times observed in an $M/M/1$ queue. In an $M/M/1$ queueing system, we let A_i denote the interarrival time between the customers numbered $i-1$ and i (with $A_0 \equiv 0$) so that $\{A_i : i = 1, 2, \ldots\} \overset{\text{i.i.d.}}{\sim} \text{Exponential}(\lambda)$ and $\mathrm{E}[A_i] = 1/\lambda$; moreover, we let B_i denote the service time of the ith customer so that $\{B_i : i = 1, 2, \ldots\} \overset{\text{i.i.d.}}{\sim} \text{Exponential}(\nu)$ and $\mathrm{E}[B_i] = 1/\nu$. If Y_i denotes the waiting time in the queue for the ith customer in this single-server queueing system, then $Y_{i+1} = \max\{0, Y_i + B_i - A_{i+1}\}$ for $i = 1, 2, \ldots$.

The $M/M/1$ queue waiting times $\{Y_i : i = 1, 2, \ldots\}$ constitute a test process with highly non-normal marginals and an autocorrelation function that decays approximately at a geometric rate. In terms of the arrival rate λ, the service rate ν, and traffic intensity $\rho = \lambda/\nu$, the process $\{Y_i\}$ has marginal distribution function

$$F_Y(y) \equiv \Pr\{Y_i \le y\} = \begin{cases} 0, & y < 0 \\ 1-\rho, & y = 0 \\ 1 - \rho e^{-(\nu-\lambda)y}, & y > 0 \end{cases} \qquad (4)$$

so that the marginal mean and variance are given by

$$\mu = \mathrm{E}[Y_i] = \frac{\rho^2}{\lambda(1-\rho)}, \quad \sigma^2 = \mathrm{Var}[Y_i] = \frac{\rho^3(2-\rho)}{\lambda^2(1-\rho)^2} \qquad (5)$$

respectively. The lag-ℓ covariance of the process $\{Y_i\}$ is

$$\text{Cov}(Y_i, Y_{i+\ell}) = \frac{1-\rho^2}{2\pi\lambda^2}\int_0^r \frac{z^{|\ell|+3/2}(r-z)^{1/2}}{(1-z)^3}\,dz \quad (6)$$

for $\ell = 0, \pm 1, \pm 2, \ldots$, where $r = 4\rho/(1+\rho)^2$ so that $0 < r < 1$; and the variance parameter is given by

$$\Omega^2 = \frac{\rho^3\left(\rho^3 - 4\rho^2 + 5\rho + 2\right)}{\lambda^2(1-\rho)^4}. \quad (7)$$

The service rate of the in-control process is set to $\nu_0 = 1$. To test different levels of dependence, we take the in-control arrival rate $\lambda_0 \in \{0.3, 0.6\}$ so that for the traffic intensity of the in-control system, we have $\rho_0 \in \{0.3, 0.6\}$. We generate the monitored process $\{Y_i : i = 1, 2, \}$ based on the algorithm of Schmeiser and Song (1989) so that the process is stationary with the steady-state properties (4)–(7).

Similar to AR(1) processes, the shift ratio Ω^2/Ω_0^2 varies over 1, 1.5, 2, 2.5, 3, 4, and 10. For out-of-control processes, we assume that there is no change in the arrival rate λ, but the shift is caused by a change in ν. To generated shifted data, for given λ_0 and shifted Ω^2, we search for the value of ν by solving (7) for ν. Batch sizes are determined by an algorithm in Kim (2006). The DFTCV chart with the batch sizes results in actual ARL_0 close to the target value and were able to detect shifts as show in Table 3.

5 CONCLUSIONS

We propose a distribution-free SPC chart for monitoring variability of autocorrelated processes by combining a distribution-free variance estimation technique with a distribution-free SPC chart for monitoring mean of autocorrelated processes. The proposed chart uses variance estimates from nonoverlapped batches as basic observations. Although we propose only one SPC chart for monitoring variability in this paper, a number of different SPC charts can be proposed by considering other variance estimation techniques (e.g., the weighted Area estimator or other improved versions of standardized time series estimators) and other SPC charts for monitoring mean (e.g., SPC charts due to Johnson and Bagshaw 1974 and Kim et al. 2006a).

The performance of the DFTCV chart could be improved if a variance estimator with better statistical properties—low bias and low variance—is used to get a variance estimate from each batch.

It will be interesting to explore the performance of the proposed SPC chart with other variance estimators. Most SPC charts in literature assume that there exist large number of training data points and that estimated values are close to the true values. This paper also assumes that parameters such as Ω^2 is known. However, a research that investigates the impact of using estimated values from a training data set of a moderate size is currently undergoing.

Table 1: Two-Sided ARLs in Terms of Number of Raw Observations for an AR(1) Process when Changes Are Caused By σ^2 or σ_ε^2 (Shift Ratios Are in the Units of Ω^2/Ω_0^2)

φ	Shift Ratio	ARL
0.25	1	11899
($m = 16$)	1.5	1037
	2	450
	2.5	293
	3	224
	4	152
	10	63
0.3	1	10704
($m = 16$)	1.5	1115
	2	475
	2.5	306
	3	231
	4	156
	10	65
0.5	1	11087
($m = 31$)	1.5	1732
	2	778
	2.5	506
	3	375
	4	258
	10	110
0.7	1	11681
($m = 60$)	1.5	2677
	2	1210
	2.5	818
	3	613
	4	427
	10	192
0.9	1	11520
($m = 166$)	1.5	5144
	2	2497
	2.5	1669
	3	1285
	4	932
	10	437

REFERENCES

Alexopoulos, C., D. Goldsman, and R. F. Serfozo. 2005. Stationary processes: statistical estimation. In *The Encyclopedia of Statistical Sciences*, ed. N. L. Johnson, and C. B. Read, 2nd edition. New York: John Wiley & Sons.

Table 2: Two-Sided ARLs in Terms of Number of Raw Observations for an AR(1) Process when Changes Are Caused by φ (Shift Ratios Are in the Units of Ω^2/Ω_0^2)

φ	Shift Ratio	ARL
0.25	1	11948
($m = 16$)	1.5	1571
	2	757
	2.5	560
	3	474
	4	418
	10	562
0.3	1	10647
($m = 16$)	1.5	1873
	2	897
	2.5	662
	3	576
	4	507
	10	851
0.5	1	11351
($m = 31$)	1.5	2461
	2	1220
	2.5	861
	3	718
	4	616
	10	737
0.7	1	11484
($m = 60$)	1.5	3598
	2	1848
	2.5	1327
	3	1117
	4	964
	10	1094
0.9	1	11648
($m = 166$)	1.5	6837
	2	4110
	2.5	3187
	3	2874
	4	2502
	10	3678

Table 3: Two-Sided ARLs in Terms of Number of Raw Observations for an M/M/1 Queue (Shift Ratios Are in the Units of Ω^2/Ω_0^2)

ρ	Shift Ratio	ARL
0.3	1	9721
($m = 659$)	1.5	6950
	2	4735
	2.5	3703
	3	3073
	4	2361
	10	1343
0.6	1	9217
($m = 931$)	1.5	7843
	2	6344
	2.5	5251
	3	4411
	4	3480
	10	2106

Billingsley, P. 1968. *Convergence of probability measures*. New York: John Wiley & Sons.

Chiu, C. C., M.-K. Chen, and K.-M. Lee. 2001. Shifts recognition in correlated process data using a neural network. *International Journal of Systems Science* 32:137–143.

Cook, D. F., and C. C. Chiu. 1998. Using radial basis function neural networks to recognize shifts in correlated manufacturing process parameters. *IIE Transactions* 30:227–234.

Cook, D. F., C. W. Zobel, and Q. J. Nottingham. 2001. Utilization of neural networks for the recognition of variance shifts in correlated manufacturing process parameters. *International Journal of Production Research* 39:3881–3887.

Goldsman, D., K. Kang, and A. F. Seila. 1999. Cramér–von Mises variance estimators for simulations. *Operations Research* 47:299–309.

Johnson, R. A., and M. Bagshaw. 1974. The effect of serial correlation on the performance of CUSUM tests. *Technometrics* 16:103–112.

Kim, S.-H. 2006. Distribution-free SPC methods for monitoring variability of autocorrelated processes. Technical Report, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, Georgia.

Kim, S.-H., C. Alexopoulos, D. Goldsman, and K.-L. Tsui. 2006a. A model-free CuSum procedure for autocorrelated processes. Technical Report, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, Georgia.

Kim, S.-H., C. Alexopoulos, K.-L. Tsui, and J. R. Wilson. 2006b. A distribution-free tabular CUSUM chart for autocorrelated data. *IIE Transactions*, forthcoming.

Lada, E. K., and J. R. Wilson. 2005. A wavelet-based spectral procedure for steady-state simulation analysis. To appear in *European Journal of Operational Research*.

Lu, C.-W., and M. R. Reynolds, Jr. 1999. Control charts for monitoring the mean and variance of autocorrelated processes. *Journal of Quality Technology* 31:259–274.

Reynolds, M. R., Jr. and C.-W. Lu. 1997. Control charts for monitoring processes with autocorrelated Data. *Non-*

linear Analysis, Theory, Methods, and Applications 30:4059–4067.

Runger, G. C. and T. R. Willemain. 1995. Model-based and model-free control of autocorrelated processes. *Journal of Quality Technology* 27:283–292.

Shapiro, S. S., and M. B. Wilk. 1965. An analysis of variance test for normality (complete samples). *Biometrika* 52:591–611.

Schmeiser, B. W., and W. T. Song. 1989. Inverse-transformation algorithms for some common stochastic processes. In *Proceedings of the 1989 Winter Simulation conference*, ed. E. A. MacNair, K. J. Musselman, and P. Heidelberger, 490–496.

Schruben, L. W. 1983. Confidence interval estimation using standardized time series. *Operations Research* 31:1090–1108.

von Neumann, J. 1941. Distribution of the ratio of the mean square successive difference to the variance. *The Annals of Mathematical Statistics* 12:367–395.

Yashchin, E. 1993. Performance of CUSUM control schemes for serially correlated observations. *Technometrics* 35:37–52.

AUTHOR BIOGRAPHIES

SEONG-HEE KIM is an Assistant Professor in the School of Industrial and Systems Engineering at the Georgia Institute of Technology. She received her Ph.D. in Industrial Engineering and Management Sciences from Northwestern University in 2001. Her research interests include optimization via simulation, quality control, and simulation output analysis. She is a member of INFORMS and IIE, and serves the editorial board of the Simulation Department of *IIE Transactions*. She is also an associate editor in the simulation area of *Operations Research* and OR/simulation area of *The American Statistician*. Her e-mail and web addresses are <skim@isye.gatech.edu> and <www.isye.gatech.edu/~skim/>, respectively.