

A SEQUENTIAL METHOD FOR ESTIMATING STEADY-STATE QUANTILES USING STANDARDIZED TIME SERIES

Athanasios Lolos
J. Haden Boone
Christos Alexopoulos
David Goldsman

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

Kemal Dinçer Dineç

Department of Industrial Engineering
Gebze Technical University
41400 Gebze, Kocaeli, TURKEY

Anup C. Mokashi

Memorial Sloan Kettering Cancer Center
1275 York Avenue
New York NY 10065, USA

James R. Wilson

Edward P. Fitts Department of Industrial and Systems Engineering
North Carolina State University
Raleigh, NC 27695-7906, USA

ABSTRACT

We propose SQSTS, an automated sequential procedure for computing confidence intervals (CIs) for steady-state quantiles based on Standardized Time Series (STS) processes computed from sample quantiles. We estimate the variance parameter associated with a given quantile estimator using the order statistics of the full sample and a combination of variance-parameter estimators based on the theoretical framework developed by Alexopoulos et al. in 2022. SQSTS is structurally less complicated than its main competitors, the Sequest and Sequem methods developed by Alexopoulos et al. in 2019 and 2017. Preliminary experimentation with the customer delay process prior to service in a congested M/M/1 queueing system revealed that SQSTS performed favorably compared with Sequest and Sequem in terms of the estimated CI coverage probability, and it significantly outperformed the latter methods with regard to average sample-size requirements.

1 INTRODUCTION

The estimation of steady-state quantiles via simulation experiments is crucial in the design and performance assessment of complex systems. For instance, marginal steady-state quantiles for cycle times play a critical role in the design of manufacturing systems (Hopp and Spearman 2008) as well as for contracts between manufacturers and clients. Further, this problem is central in financial engineering applications (Glasserman 2004). However, the growth in this research area has been slow during the last few decades because of the theoretical and computational challenges associated with quantile estimation such as: (i) contamination by initialization bias; (ii) substantial autocorrelation between successive outputs; and (iii) a variety of issues

associated with the marginal probability density function (p.d.f.) including discontinuities, multimodalities with sharp peaks (Alexopoulos et al. 2018), and several departures from global smoothness. As a result, the list of two-stage and sequential methods for steady-state quantile estimation is relatively limited compared with the numerous sequential procedures for estimating the steady-state mean. Alexopoulos et al. (2019) give an extensive review of earlier methods including Raatikainen (1990) and Chen and Kelton (2006, 2008). Notably, the Sequest method in the former paper outperformed the two-phase QI procedure of Chen and Kelton (2008) with regard to sampling efficiency. With regard to the estimation of the steady-state mean, Dong and Glynn (2019) lay a foundation for sequential STS-based procedures that yield asymptotically exact CIs. Their set of sufficient conditions includes the strong approximation assumption of Damerdjji (1994); certain regularity conditions involving the behavior of the sequential procedure as a function of the current simulation clock and sample path (Dong and Glynn 2019, p. 334); and the weak convergence, as the upper bound on the CI's precision approaches zero, of the denominator of the respective pivotal random variable (r.v.) at the procedure's stopping time to a r.v. that is positive almost surely (a.s.). To the best of our knowledge, the last condition has not been formally proven when the number of batches exceeds 3.

In this article we present the first sequential procedure for estimating steady-state quantiles based on STS processes that are computed from nonoverlapping batches. Our SQSTS method draws elements from two recent procedures with different objectives: the aforementioned Sequest procedure for quantile estimation and the SPSTS method of Alexopoulos et al. (2016) for the estimation of the steady-state mean. We temporarily bypass the Sequem procedure of Alexopoulos et al. (2017) because it is an extension of Sequest for extreme quantiles ($p \geq 0.95$ or $p \leq 0.05$) based on the maximum-transformation technique (Heidelberger and Lewis 1984). Although the STS-based estimation of the mean dates back to the 1980s (Schruben 1983; Glynn and Iglehart 1990; Goldsman et al. 1990), the use of STS for quantile estimation is a rather recent development. Calvin and Nakayama (2013) proposed this methodology for independent and identically distributed (i.i.d.) data. Subsequently Alexopoulos et al. (2020, 2022) laid out the theoretical foundations for STS-based quantile estimation in dependent processes, established asymptotic properties for a variety of variance-parameter estimators for the sample quantile computed from nonoverlapping batches, and closed various theoretical gaps from the 1980s related to STS-based variance-parameter estimation. In particular, a variance-parameter estimator computed from a linear combination of the average of batched STS area estimators and a modified sample variance of the sample quantiles from the same batches converges to a scaled chi-squared random variable with nearly twice the degrees of freedom (d.f.) than each of its constituents as the batch size tends to infinity while the number of batches is held constant.

Compared with the Sequest and SPSTS procedures, the proposed method has the following fundamental differences: (i) it is substantially simpler than Sequest (see Section 3 for details); (ii) it alters the structure of SPSTS with modifications aiming to address issues pertinent to small-sample bias of the STS-based variance-parameter estimator; (iii) it overcomes an ad hoc compensation for the variance-parameter estimator used in SPSTS to resolve small-sample-bias issues; and (iv) it is statistically more efficient than Sequest because of the adoption of the combined variance-parameter estimator.

Section 2 contains important background information and reviews the core assumptions and theorems that form the basis of our sequential method. Section 3 contains a description of the SQSTS algorithm, and in Section 4 the performance of SQSTS is tested against the recent Sequest and Sequem procedures. Finally, Section 5 summarizes our work and discusses future directions.

2 FOUNDATIONS

This section discusses the basic notation, the assumptions, and the core results that constitute the foundations of the SQSTS sequential procedure.

2.1 Notation

For $p \in (0, 1)$, the p -quantile of a r.v. X is the inverse of the cumulative distribution function (c.d.f.) $F(x)$, $x_p \equiv F^{-1}(p) \equiv \inf\{x : F(x) \geq p\}$. The primary goal is to compute a point estimate and a CI for x_p based on a simulation-generated finite sample $\{X_1, X_2, \dots, X_n\}$ of size $n \geq 1$. The estimation of x_p is based on the stationary time series $\{Y_k : k \geq 1\}$, which is a warmed-up (i.e., truncated and reindexed) version of the original sequence of simulation output $\{X_i : i \geq 1\}$. Let $Y_{(1)} \leq \dots \leq Y_{(n)}$ be the respective order statistics. The classical point estimator of x_p is the empirical p -quantile $\tilde{y}_p(n) \equiv Y_{(\lceil np \rceil)}$, where $\lceil \cdot \rceil$ denotes the ceiling function; and for completeness we take $Y_{(0)} \equiv 0$ to handle anomalous conditions.

For each $x \in \mathbb{R}$ and $k \geq 1$, we define the indicator r.v. as $I_k(x) \equiv 1$ if $Y_k \leq x$, and $I_k(x) \equiv 0$ otherwise; hence $E[I_k(x_p)] = p$. Assuming $n \geq 2$, we let $\bar{I}(x_p, n) \equiv n^{-1} \sum_{k=1}^n I_k(x_p)$; and for each $\ell \in \mathbb{Z}$, we let $\rho_I(x_p, \ell) \equiv \text{Corr}[I_k(x_p), I_{k+\ell}(x_p)]$ denote the autocorrelation function of the indicator process $\{I_k(x_p) : k \geq 1\}$ at lag ℓ . Below we also adopt the following notation: $N(0, 1)$ denotes the standard normal distribution; $\mathbf{Z}_\nu \equiv [Z_1, \dots, Z_\nu]$ denotes a ν -dimensional vector whose components are i.i.d. $N(0, 1)$; χ_ν^2 denotes a chi-squared r.v. with ν d.f.; t_ν denotes a r.v. having Student's t distribution with ν d.f.; and $t_{\delta, \nu}$ denotes the δ -quantile of t_ν . We let D denote the space of real-valued functions on $[0, 1]$ that are right-continuous with left-hand limits (Whitt 2002, §3.3).

The assumptions and the core results that are outlined below are the key elements for variance cancellation methods to develop asymptotically exact $100(1 - \alpha)\%$ CIs for x_p with form $\tilde{y}_p(n) \pm t_{1-\alpha/2, \nu} \hat{\sigma} / \sqrt{n}$, where $\hat{\sigma}^2$ is the estimator of the (quantile) variance parameter $\sigma^2 \equiv \lim_{n \rightarrow \infty} n \text{Var}[\tilde{y}_p(n)]$ and the d.f. ν depends on the underlying method.

2.2 Assumptions

This section contains the key assumptions for the processes $\{Y_k : k \geq 1\}$ and $\{I_k(x_p) : k \geq 1\}$.

Geometric-Moment Contraction (GMC) Condition (Wu 2005). The process $\{Y_k : k \geq 1\}$ is defined by a function $\xi(\cdot)$ of a sequence of i.i.d. random variables $\{\varepsilon_j : j \in \mathbb{Z}\}$ such that $Y_k = \xi(\dots, \varepsilon_{k-1}, \varepsilon_k)$ for $k \geq 0$. Moreover, there exist constants $\psi > 0$, $C > 0$, and $r \in (0, 1)$ such that for two independent sequences $\{\varepsilon_j : j \in \mathbb{Z}\}$ and $\{\varepsilon'_j : j \in \mathbb{Z}\}$ each consisting of i.i.d. variables distributed like ε_0 , we have

$$E[|\xi(\dots, \varepsilon_{-1}, \varepsilon_0, \varepsilon_1, \dots, \varepsilon_k) - \xi(\dots, \varepsilon'_{-1}, \varepsilon'_0, \varepsilon_1, \dots, \varepsilon_k)|^\psi] \leq Cr^k \text{ for } k \geq 0.$$

The GMC condition holds for a plethora of processes including autoregressive moving-average time series (Shao and Wu 2007), a rich collection of processes with short-range dependence, and a broad class of Markov chains; see Alexopoulos et al. (2019, 2022) for an extended list of citations and empirical methods for verifying the GMC assumption in practice. Recently, we have established the validity of the GMC condition for the customer delay (before service) process in an M/M/1 queueing system and a G/G/1 system with non-heavy-tailed service-time distributions (Dineç et al. 2022b).

Density-Regularity (DR) Condition. The p.d.f. $f(\cdot)$ is bounded on \mathbb{R} and continuous almost everywhere (a.e.) on \mathbb{R} ; moreover, $f(x_p) > 0$, and the derivative $f'(x_p)$ exists.

Short-Range Dependence (SRD) of the Indicator Process. The indicator process $\{I_k(x_p) : k \geq 1\}$ has the SRD property so that

$$0 < \sum_{\ell \in \mathbb{Z}} \rho_I(x_p, \ell) \leq \sum_{\ell \in \mathbb{Z}} |\rho_I(x_p, \ell)| < \infty. \quad (1)$$

Thus the variance parameters for the r.v.'s $\bar{I}(x_p, n)$ and $\tilde{y}_p(n)$ satisfy the relations

$$\left. \begin{aligned} \sigma_I^2 &\equiv \lim_{n \rightarrow \infty} n \text{Var}[\bar{I}(x_p, n)] = p(1-p) \sum_{\ell \in \mathbb{Z}} \rho_I(x_p, \ell) \in (0, \infty), \\ \sigma^2 &= \lim_{n \rightarrow \infty} n \text{Var}[\tilde{y}_p(n)] = \frac{\sigma_I^2}{f^2(x_p)} \in (0, \infty). \end{aligned} \right\} \quad (2)$$

Functional Central Limit Theorem (FCLT) for the Indicator Process. We define the following sequence of random functions $\{\mathcal{J}_n : n \geq 1\}$ in D ,

$$\mathcal{J}_n(t; x_p) \equiv \frac{\lfloor nt \rfloor}{\sigma_I n^{1/2}} [\bar{I}(x_p, \lfloor nt \rfloor) - p] \quad \text{for } t \in [0, 1] \text{ and } n \geq 1,$$

where $\lfloor \cdot \rfloor$ denotes the floor function. We assume that this random-function sequence satisfies the FCLT

$$\mathcal{J}_n \xrightarrow[n \rightarrow \infty]{\Rightarrow} \mathcal{W}$$

in D with the appropriate metric, where \mathcal{W} denotes standard Brownian motion on $[0, 1]$; and $\xrightarrow[n \rightarrow \infty]{\Rightarrow}$ denotes weak convergence as $n \rightarrow \infty$ (Billingsley 1999, pp. 1–6 and Theorem 2.1).

Remark 1 If the SRD condition holds, then to all intents and purposes it is reasonable to assume the validity of the FCLT above (Whitt 2002, p. 107). Recently Dengeç et al. (2022a) proved that if the stationary process $\{Y_k : k \geq 1\}$ obeys the GMC condition, then the associated indicator process $\{I_k(x_p) : k \geq 1\}$ satisfies the SRD condition. This result provides solid theoretical and practical evidence of the mutual compatibility of the GMC, SRD, and FCLT conditions.

2.3 Asymptotic Properties Based on Nonoverlapping Batches

The SQSTS sequential procedure is based on nonoverlapping batches. Given a fixed batch count $b \geq 2$, for $j = 1, \dots, b$ the j th nonoverlapping batch of size $m \geq 1$ consists of the subsequence $\{Y_{(j-1)m+1}, \dots, Y_{jm}\}$, where we assume $n = bm$. The batch mean of the associated indicator r.v.'s for the j th batch is $\bar{I}_j(x_p, m) \equiv m^{-1} \sum_{\ell=1}^m I_{(j-1)m+\ell}(x_p)$. Similar to the full-sample case, we define the order statistics $Y_{j,(1)} \leq \dots \leq Y_{j,(m)}$ corresponding to the j th batch. Then the j th batched quantile estimator (BQE) of x_p is $\hat{y}_p(j, m) \equiv Y_{j,(\lfloor mp \rfloor)}$.

Theorem 1 (Alexopoulos et al. 2019) If the output process $\{Y_k : k \geq 1\}$ satisfies the GMC and DR conditions, and the indicator process $\{I_k(x_p) : k \geq 1\}$ satisfies the SRD and the respective FCLT conditions, then we obtain the Bahadur representation

$$\hat{y}_p(j, m) = x_p - \frac{\bar{I}_j(x_p, m) - p}{f(x_p)} + O_{\text{a.s.}} \left[\frac{(\log m)^{3/2}}{m^{3/4}} \right]$$

as $m \rightarrow \infty$ for $j = 1, \dots, b$, where the big- $O_{\text{a.s.}}$ notation for the remainder,

$$Q_{j,m} \equiv \hat{y}_p(j, m) - x_p + \frac{\bar{I}_j(x_p, m) - p}{f(x_p)},$$

means that there exist r.v.'s \mathcal{U}_j and \mathcal{R}_j that are bounded a.s. and satisfy

$$|Q_{j,m}| \leq \mathcal{U}_j \frac{(\log m)^{3/2}}{m^{3/4}} \quad \text{for } m \geq \mathcal{R}_j \quad \text{a.s.}$$

Further,

$$m^{1/2} [\hat{y}_p(1, m) - x_p, \dots, \hat{y}_p(b, m) - x_p] \xrightarrow[m \rightarrow \infty]{\Rightarrow} \sigma \mathbf{Z}_b$$

in \mathbb{R}^b with the standard Euclidean metric.

2.4 Standardized Time Series for Quantile Estimation

The full-sample STS for quantile estimation is defined as

$$T_n(t) \equiv \frac{\lfloor nt \rfloor}{n^{1/2}} [\tilde{y}_p(n) - \tilde{y}_p(\lfloor nt \rfloor)] \quad \text{for } n \geq 1 \text{ and } t \in [0, 1].$$

We have the following key result.

Theorem 2 (Alexopoulos et al. 2022) If $\{Y_k : k \geq 1\}$ satisfies the assumptions of Theorem 1, then in $\mathbb{R} \times D$,

$$[n^{1/2}(\tilde{y}_p(n) - x_p), T_n] \xrightarrow[n \rightarrow \infty]{\Rightarrow} \sigma[\mathcal{W}(1), \mathcal{B}],$$

where $\mathcal{B}(\cdot)$ is a standard Brownian bridge process that is independent of $\mathcal{W}(1)$.

The full-sample STS area estimator of the variance parameter σ^2 is $A_n^2(w)$, where

$$A_n(w) \equiv n^{-1} \sum_{k=1}^n w(k/n) T_n(k/n) \quad \text{for } n \geq 1,$$

and $w(\cdot)$ is a deterministic function that is bounded and continuous a.e. in $[0, 1]$ such that the r.v. $Z(w) \equiv \int_0^1 w(t) \mathcal{B}(t) dt \sim N(0, 1)$.

Some of the weight functions that satisfy the above conditions are the constant $w_0(t) = \sqrt{12}$ (Schruben 1983), the quadratic $w_2(t) = \sqrt{840}(3t^2 - 3t + 1/2)$ (Goldsman et al. 1990), and the orthonormal sequence $w_{\cos,j}(t) = \sqrt{8}\pi j \cos(2\pi jt)$, $j = 1, 2, \dots$ (Foley and Goldsman 1999). The latter two classes yield first-order unbiased estimators for the variance parameter $\lim_{n \rightarrow \infty} n \text{Var}(\bar{Y}_n)$ related to the sample mean $\bar{Y}_n \equiv n^{-1} \sum_{k=1}^n Y_k$ of the base process $\{Y_k : k \geq 1\}$; hence they were tailored to the estimation of the steady-state mean. In our experiment in Section 4 we used only the constant weight w_0 because preliminary experimentation with the other alternatives on two test processes, including the process in Section 4, did not reveal substantive benefits with regard to the small-sample bias of the area estimator $A_n^2(w)$.

Theorem 3 (Alexopoulos et al. 2022) If $\{Y_k : k \geq 1\}$ satisfies the assumptions of Theorem 1, then

$$A_n^2(w) \xrightarrow[n \rightarrow \infty]{\Rightarrow} \sigma^2 \chi_1^2.$$

These results can be extended for the case of nonoverlapping batches. Specifically, we let

$$T_{j,m}(t) \equiv \frac{\lfloor mt \rfloor}{m^{1/2}} [\hat{y}_p(j, m) - \hat{y}_p(j, \lfloor mt \rfloor)] \quad \text{for } m \geq 1, \quad 1 \leq j \leq b, \quad \text{and } t \in [0, 1]$$

be the STS-based quantile-estimation process for the j th batch, where

$$\hat{y}_p(j, \lfloor mt \rfloor) \equiv \left\{ \begin{array}{ll} 0, & \text{if } \lfloor mt \rfloor = 0; \\ \text{the empirical } p\text{-quantile (i.e., the } \lceil p \lfloor mt \rceil \rceil\text{-th order statistic)} & \\ \text{computed from the partial sample } \{Y_{(j-1)m+k} : k = 1, \dots, \lfloor mt \rfloor\}, & \text{otherwise.} \end{array} \right\} \quad (3)$$

We define the signed area computed from batch j as

$$A_{j,m}(w) \equiv m^{-1} \sum_{k=1}^m w(k/m) T_{j,m}(k/m) \quad \text{for } j = 1, \dots, b. \quad (4)$$

The batched STS-area estimator is

$$\mathcal{A}_{b,m}^2(w) \equiv b^{-1} \sum_{j=1}^b A_{j,m}^2(w). \quad (5)$$

Theorems 4 and 5 below outline the asymptotic validity of the CIs used in SQSTS and the Sequest method of Alexopoulos et al. (2019).

Theorem 4 (Alexopoulos et al. 2022) If $\{Y_k : k \geq 1\}$ satisfies the assumptions of Theorem 1, then the vector of the signed areas $[A_{1,m}(w), \dots, A_{b,m}(w)]$ converges weakly to the same distributional limit as the scaled vector of BQEs in Theorem 1:

$$[A_{1,m}(w), \dots, A_{b,m}(w)] \xrightarrow{m \rightarrow \infty} \sigma \mathbf{Z}_b.$$

Further,

$$\mathcal{A}_{b,m}^2(w) \xrightarrow{m \rightarrow \infty} \sigma^2 \chi_b^2/b, \quad (6)$$

and

$$\tilde{y}_p(n) \pm t_{1-\alpha/2,b}(\mathcal{A}_{b,m}^2(w)/n)^{1/2} \quad (7)$$

is an asymptotically exact $100(1 - \alpha)\%$ CI for x_p .

Finally, we define the mean squared deviation of the BQEs from the full-sample estimator $\tilde{y}_p(n)$,

$$\tilde{S}_{b,m}^2 \equiv (b-1)^{-1} \sum_{j=1}^b [\hat{y}_p(j,m) - \tilde{y}_p(n)]^2, \quad (8)$$

and the combined variance-parameter estimator

$$\tilde{\mathcal{V}}_{b,m}(w) \equiv \frac{b\mathcal{A}_{b,m}^2(w) + (b-1)m\tilde{S}_{b,m}^2}{2b-1}.$$

Theorem 5 (Alexopoulos et al. 2022) If $\{Y_k : k \geq 1\}$ satisfies the assumptions of Theorem 1, then

$$n^{1/2}[\tilde{y}_p(n) - x_p] \xrightarrow{m \rightarrow \infty} \sigma \mathbf{Z}_1, \quad (9)$$

$$m\tilde{S}_{b,m}^2 \xrightarrow{m \rightarrow \infty} \sigma^2 \chi_{b-1}^2/(b-1), \quad (10)$$

$$\tilde{\mathcal{V}}_{b,m}(w) \xrightarrow{m \rightarrow \infty} \sigma^2 \chi_{2b-1}^2/(2b-1), \quad (11)$$

the limiting r.v.'s in Equations (6), (9), and (10) are independent, and the limiting r.v.'s in Equations (9), and (11) are also independent. Further,

$$\tilde{y}_p(n) \pm t_{1-\alpha/2,b-1}(m\tilde{S}_{b,m}^2/n)^{1/2} \quad (12)$$

and

$$\tilde{y}_p(n) \pm t_{1-\alpha/2,2b-1}(\tilde{\mathcal{V}}_{b,m}(w)/n)^{1/2} \quad (13)$$

are also asymptotically exact $100(1 - \alpha)\%$ CI estimators of x_p .

Hereafter, we refer to $m\tilde{S}_{b,m}^2$ as the nonoverlapping batched quantile (NBQ) variance-parameter estimator. The CI in Equation (12) has been used in the Sequest method (Alexopoulos et al. 2019). The benefits of the combined variance-parameter estimator $\tilde{\mathcal{V}}_{b,m}(w)$ are apparent: since its distributional limit as $m \rightarrow \infty$ has nearly double d.f. than its constituents $\mathcal{A}_{b,m}^2(w)$ and $m\tilde{S}_{b,m}^2$ for large m , the CI in Equation (13) will have significantly less variable half-length (by a factor of about $\sqrt{2}$) than each of the two competitors in Equations (7) and (12); this typically results in better sampling efficiency. Limited experimentation in

Alexopoulos et al. (2022) has revealed that the batched area estimator $\mathcal{A}_{b,m}^2(w_0)$ based on the constant weight function $w_0(t) = \sqrt{12}$ is substantially more biased for small batch sizes m than its counterpart $m\tilde{S}_{b,m}^2$ based on the BQEs. This small-sample-bias problem for STS-based estimators has been known since the 1980s (relative to the estimation of the steady-state mean), but it appears to be more pronounced with regard to quantile estimation. Although the combined estimator $\tilde{\mathcal{V}}_{b,m}(w)$ partially rectifies this problem, based on the limited experimentation in Alexopoulos et al. (2022), SQSTS takes more aggressive initial steps to remove excessive bias from the batched area estimator $\mathcal{A}_{b,m}^2(w)$ that is due to small batch sizes.

On the computational front, the variance-parameter estimator $m\tilde{S}_{b,m}^2$ can be computed in $O(n \log_2 n)$ time using a quicksort algorithm for (primitive) arrays. On the other hand, the computation of the area estimator $\mathcal{A}_{b,m}^2(w)$ (and the combined estimator $\tilde{\mathcal{V}}_{b,m}(w)$) has the same worst-case complexity if one uses a more-complex data structure with objects, at the cost of higher memory usage (Alexopoulos et al. 2022).

3 SEQUENTIAL PROCEDURE

In this section we present our sequential procedure for estimating steady-state quantiles of a stationary process $\{Y_k : k \geq 1\}$. The core of the SQSTS procedure consists of three loops. The first loop progressively increases the batch size m until the signed areas $A_{j,m}(w)$ pass the two-sided randomness test of von Neumann (1941), while the second loop increases the batch size until the signed areas pass the one-sided test of Shapiro and Wilk (1965) for testing the hypothesis that the (nearly) i.i.d. sample $\{A_{j,m}(w) : j = 1, \dots, b\}$ has a univariate normal distribution with unspecified mean and standard deviation. To control the growth of the batch size, both loops use a rapidly decreasing sequence of significance levels. We focus on the signed areas in an attempt to mitigate the issues caused by the pronounced small-batch bias of the batched area estimator $\mathcal{A}_{b,m}^2(w)$ compared to the NBQ variance-parameter estimator. As we mentioned earlier, this assessment is based only on limited experimentation with two test processes, including the process in Section 4. At the end of the two loops, the signed areas $A_{j,m}(w)$ satisfy approximately the asymptotic properties in Theorems 4–5. The last loop of SQSTS starts with a rebatch of the current time series that quadruples the batch size and then performs iterative increases of the batch count b or batch size m until the CI for x_p in Equation (13) meets the target relative-precision requirement, provided such a requirement has been specified. The next two paragraphs provide a brief description of each step of SQSTS.

Steps [0]–[1] initialize the experimental parameters and generate the initial data comprised of $b = 64$ batches of size 512 when $p \in [0.05, 0.95]$ or 4096 otherwise. The user inputs the value of α (CI error probability) and, potentially, an upper bound r^* on the CI relative precision. The level of significance for the statistical tests in Steps [2]–[3] is set according to the sequence $\{\beta\psi(\ell) : \ell = 1, 2, \dots\}$, where $\beta = 0.3$, $\psi(\ell) \equiv \exp[-\eta(\ell - 1)^\theta]$, $\eta = 0.2$, and $\theta = 2.3$. Step [2] consists of a loop that assesses the extent to which the signed areas $A_{j,m}(w)$ are (nearly) i.i.d. using the two-sided test of von Neumann with progressively decreasing size $\beta\psi(\ell)$ on iteration ℓ . The large initial values of the batch count b , the type-I error $\beta\psi(1)$, and the batch size m aim at increasing the power of von Neumann’s test. For instance, the (normal) null distribution of the latter test can be badly distorted by departures from normality in the signed areas $A_{j,m}(w)$; this forms the basis of starting with a relatively large batch size. The values η and θ were chosen to avoid excessive incremental increases in the batch size. Notice that on iteration $\ell = 4$ one has $\beta\psi(4) = 0.025$. If the signed areas fail the randomness test, the batch size is incremented by the factor of $\sqrt{2}$ and $b(\lceil m\sqrt{2} \rceil - m)$ additional data are generated, where the function $\lceil \cdot \rceil$ rounds its argument to the nearest integer. At the end of Step [2], the signed areas are nearly i.i.d. Step [3] contains a second loop that assesses the univariate normality of the signed areas $A_{j,m}(w)$ using the one-sided Shapiro-Wilk test, again with level of significance $\beta\psi(\ell)$ on iteration ℓ . It should be noted that the Shapiro-Wilk test is widely recognized for having the highest power among several alternative tests for univariate normality. In particular, it is most powerful when the data have a continuous, skewed, and short- or long-tailed distribution (Fishman 1978, Section 2.10).

Step [4] deals with the initial transient phase. Specifically, after the signed areas pass both the randomness and normality tests, the first of the 64 batches is removed and a new batch is generated in anticipation that once the latter statistical tests are passed any transient effects are restricted to the first batch. We realize that this truncation may be excessive, and plan to address it in the future. Step [5] rebatches the current sample into 16 batches of quadruple batch size. This step is designed primarily for the case where there is no precision requirement for the CI half-length; this is typical for most commercial simulation packages and a reasonable starting point for estimating the sample size required to achieve a given precision requirement. If the user has specified a relative CI precision requirement r^* , Step [6] sequentially increases the batch count b (up to $b^* = 64$) or the batch size m until the half-length of the CI for x_p does not exceed $r^*|\widehat{y}_p(n)|$. The function $\text{mid}(\cdot)$ denotes the median of its arguments. Notice that the potential increases to the batch size are constrained between 5% and 30%. Step [7] delivers the final CI for x_p based on Equation (13).

In comparison with the Sequest and Sequem procedures of Alexopoulos et al. (2017, 2019), the SQSTS algorithm is structurally less complicated. For example: (i) while Sequest starts with a smaller initial batch size (128 versus 512 or 4096), it contains an intricate loop that increases the batch size in a progressively cautious fashion until the estimated absolute skewness of the BQEs $\{\widehat{y}_p(j, m) : j = 1, \dots, b\}$ drops below an upper bound that is a function of p ; (ii) the CI for x_p delivered by Sequest incorporates adjustments for residual skewness and autocorrelation in the BQEs; and (iii) Sequem adds more complexity to Sequest because it uses groups of batches in order to apply the maximum transformation. On another front, whereas SQSTS has similar core logic to the SPSTS procedure of Alexopoulos et al. (2016) for estimating the steady-state mean, it has key differences from the latter. For instance, SPSTS attempts to control the excessive small-sample bias of the STS-based estimates of the associated variance parameter $\lim_{n \rightarrow \infty} n \text{Var}[\bar{Y}_n]$ by means of an ad hoc variance-parameter estimator computed as the maximum of the area estimators based on the cosine weights $w_{\cos,1}(\cdot)$ and $w_{\cos,2}(\cdot)$ and corresponding variance-parameter estimator resulting from the method of overlapping batch means (Meketon and Schmeiser 1984). SQSTS provides an additional safeguard against small-sample bias with the aggressive rebatching in Step [5]. At this juncture we are compelled to admit that the various heuristic assignments above severely affect our ability to establish the asymptotic validity of SQSTS. Instead, we plan to evaluate it thoroughly via an extensive experimental test bed analogous to the one in Alexopoulos et al. (2019).

An algorithmic description of SQSTS follows. For the sake of brevity, we consider only a relative CI precision requirement (if any).

Algorithm SQSTS

-
- [0] Initialization: Read the user's input for $\alpha \in (0, 1)$ and relative CI precision requirement r^* (if any). Set $\beta = 0.30$ and $b^* = 64$. Let $w(t) = \sqrt{12}$, $t \in [0, 1]$, be the weight function, and define the error function for the hypothesis tests as $\beta\psi(\ell) \equiv \exp[-\eta(\ell - 1)^\theta]$, $\ell = 1, 2, \dots$, with $\eta = 0.2$ and $\theta = 2.3$.
 - [1] Generate $b = 64$ batches of size $m = 512$ for $p \in [0.05, 0.95]$ or 4096 otherwise. Let $\ell = 1$.
 - [2] **Until** von Neumann's test fails to reject randomness:
 - compute the signed areas $A_{j,m}(w)$ for $j = 1, \dots, b$;
 - assess the randomness of $A_{j,m}(w)$ for $j = 1, \dots, b$ using von Neumann's two-sided test at significance level $\beta\psi(\ell)$;
 - set $\ell \leftarrow \ell + 1$, generate $b(\lceil m\sqrt{2} \rceil - m)$ additional observations, and set $m \leftarrow \lceil m\sqrt{2} \rceil$.
 - End**
 - [3] Reset $\ell \leftarrow 1$.
Until the Shapiro-Wilk test fails to reject normality:
 - compute the signed areas $A_{j,m}(w)$ for $j = 1, \dots, b$;

- assess the normality of $A_{j,m}(w)$ for $j = 1, \dots, b$ using the Shapiro-Wilk one-sided test for normality at significance level $\beta\psi(\ell)$;
- set $\ell \leftarrow \ell + 1$, generate $b(\lceil m\sqrt{2} \rceil - m)$ additional observations, and set $m \leftarrow \lceil m\sqrt{2} \rceil$.

End

- [4] Remove the first batch and append a new batch of size m .
- [5] Rebatch the data with $b \leftarrow b/4 = 16$ and batches of size $m \leftarrow 4m$. Compute the updated point estimate $\tilde{y}_p(n)$ and variance-parameter estimate $\tilde{\mathcal{V}}_{b,m}(w)$. If the user has not specified an upper bound r^* on the CI relative precision, go to Step [7].
- [6] **Until** the relative half-length $h(b, m, \alpha)/|\tilde{y}_p(n)| \leq r^*$, where $h(b, m, \alpha) = t_{1-\alpha/2, 2b-1}(\tilde{\mathcal{V}}_{b,m}(w)/n)^{1/2}$:
- Compute the CI midpoint $\tilde{y}_p(n)$ and the half-length $h(b, m, \alpha)$ using the combined variance-parameter estimator $\tilde{\mathcal{V}}_{b,m}(w)$.
 - Estimate the number of batches of the current size required to satisfy the precision requirement,

$$b' = \left\lceil b \left\{ \frac{h(b, m, \alpha)}{r^* \tilde{y}_p(n)} \right\}^2 \right\rceil;$$

- Update the batch count b , the batch size m , and the total sample size n as follows:

$$\begin{aligned} b &\leftarrow \min\{b', b^*\}, \\ m &\leftarrow \begin{cases} m & \text{if } b = b', \\ \lceil m \times \text{mid}\{1.05, (b'/b), 1.3\} \rceil & \text{if } b < b', \end{cases} \\ n &\leftarrow bm; \end{aligned}$$

- Generate the necessary additional data.

End

- [7] Deliver the $100(1 - \alpha)\%$ CI: $\tilde{y}_p(n) \pm t_{1-\alpha/2, 2b-1}(\tilde{\mathcal{V}}_{b,m}(w)/n)^{1/2}$.

4 EXPERIMENTAL RESULTS

This section contains a precursory empirical study designed to test the performance of SQSTS. The comparisons are made against the Sequest (Alexopoulos et al. 2019) and Sequem (Alexopoulos et al. 2017) methods, which have undergone substantial experimental evaluation. As we mentioned earlier, we used only the constant weight function w_0 . Our test process consists of the entity-delay sequence (prior to service) in an M/M/1 queueing system with arrival rate $\lambda = 0.9$, service rate $\omega = 1$ (traffic intensity $\rho = 0.9$), and FIFO service discipline. To assess the ability of the heuristic approach in Step [4] that removes the first batch after completion of the loops in Steps [2]–[3], we initialized the system with one entity in service and 112 entities in queue. The steady-state probability of this initial state is $(1 - \rho)\rho^{113} \approx 6.752 \times 10^{-7}$, implying a high probability of a prolonged transient phase. The pronounced autocorrelation function in steady-state (with lag-200 autocorrelation for the base process near 0.30) has made this process a gold-standard testbed case for steady-state simulation analysis methods. For this process Sequest and Sequem outperformed their earlier competitors with regard to sampling efficiency, but required substantial average sample sizes to deliver reliable CIs for quantiles with $p \geq 0.9$ (even in the absence of a CI precision requirement).

Table 1 contains experimental results for SQSTS, Sequest, and Sequem at the 95% confidence level. All estimates are based on 1000 independent replications; the results for Sequest (in bold typeface) are from Table 3 of Alexopoulos et al. (2019), whereas the results for Sequem (in italic typeface) are from Table 1 of Alexopoulos et al. (2017) and are limited to values of $p \geq 0.95$. In the latter case, we do not list average batch sizes because batches are used to form $\lfloor \ln(0.9)/\ln(p) \rfloor$ groups. We selected two levels of CI

relative precision, no CI precision requirement and $r^* = 0.02$ (2% CI relative precision). Column 1 lists the set of probabilities $p \in \{0.3, 0.5, 0.7, 0.9, 0.95, 0.99, 0.995\}$, column 2 contains the respective quantiles x_p computed by c.d.f. inversion, and column 3 reports the average absolute bias of the point estimates $\tilde{y}_p(n)$. Columns 4–6 list the average 95% CI half-length (HL), the average 95% CI relative precision (\bar{r}), and the estimated coverage probability of the 95% CIs for x_p ; the standard error of the estimates in column 6 is approximately 0.007. Finally, columns 7 and 8 display the average batch size (\bar{m}) and sample size (\bar{n}), respectively.

A close examination of Table 1 reveals that, in this test problem, SQSTS substantially outperforms its two recent competitors: while all methods deliver CIs with estimated coverage probabilities near the nominal value of 0.95, with the exception of Sequest for $p > 0.95$ in the absence of a CI precision requirement, SQSTS requires substantially smaller sample sizes. For example, under no CI precision requirement and for $p = 0.95$, Sequest required a factor of $9,809,640/378,815 = 25.9$ more samples on average than SQSTS. The sample size reduction is less pronounced for $p \leq 0.7$, but remained significant. Under the stringent 2% CI relative precision requirement, the ratio of the average sample sizes reflects the smaller asymptotic variance of the combined variance-parameter estimator $\tilde{\mathcal{V}}_{b,m}(w)$.

Table 1: Experimental results for SQSTS, Sequest (in bold typeface), and Sequem (in italic typeface) of x_p for the M/M/1 process with traffic intensity 0.9 based on 1000 independent replications.

p	x_p	Avg. Bias	Avg. 95% CI HL	Avg. 95% rel. prec.	Avg. 95% \bar{r} (%)	Avg. 95% CI cov. (%)	\bar{m}	\bar{n}
No CI prec. req.								
0.3	2.513	0.055	0.150	5.974	96.3	96.3	37,483	609,093
		0.034	0.095	3.801	96.6	96.6	56,354	1,806,090
0.5	5.878	0.124	0.348	5.901	96.0	96.0	30,694	498,777
		0.007	0.185	3.149	96.6	96.6	64,229	2,058,446
0.7	10.986	0.291	0.808	7.277	96.0	96.0	27,231	442,498
		0.111	0.311	2.839	96.0	96.0	81,992	2,627,562
0.9	21.972	0.717	1.948	8.827	95.3	95.3	22,018	357,785
		0.204	0.527	2.400	96.0	96.0	183,093	5,864,109
0.95	28.904	1.031	2.634	9.088	93.7	93.7	23,312	378,815
		0.274	0.654	2.268	95.0	95.0	306,385	9,809,640
		<i>0.583</i>	<i>1.543</i>	<i>5.362</i>	<i>94.2</i>	<i>94.2</i>		<i>2,961,218</i>
0.99	44.998	0.983	2.472	5.498	93.8	93.8	152,099	2,471,614
		0.777	1.055	2.371	90.0	90.0	1,008,926	32,290,677
		<i>0.684</i>	<i>1.729</i>	<i>3.846</i>	<i>95.1</i>	<i>95.1</i>		<i>15,027,284</i>
0.995	51.930	1.262	3.128	6.027	92.7	92.7	176,113	2,861,834
		1.322	1.357	2.666	86.0	86.0	1,467,551	46,966,504
		<i>0.718</i>	<i>1.781</i>	<i>3.435</i>	<i>95.1</i>	<i>95.1</i>		<i>29,584,593</i>
CI prec. req. $r^* = 2\%$								
0.3	2.513	0.020	0.048	1.896	95.1	95.1	71,132	4,528,399
		0.017	0.045	1.777	95.6	95.6	186,504	5,970,862
0.5	5.878	0.047	0.111	1.893	94.6	94.6	56,470	3,576,460
		0.039	0.105	1.783	95.6	95.6	148,044	4,740,512
0.7	10.986	0.087	0.208	1.891	94.6	94.6	58,612	3,731,135
		0.075	0.194	1.768	95.8	95.8	156,768	5,020,393
0.9	21.972	0.169	0.416	1.893	94.6	94.6	85,310	5,461,971
		0.146	0.377	1.717	95.1	95.1	257,961	8,259,880
0.95	28.904	0.226	0.547	1.892	94.1	94.1	117,098	7,500,116
		0.184	0.483	1.671	95.9	95.9	384,836	12,320,089
		<i>0.205</i>	<i>0.520</i>	<i>1.801</i>	<i>96.1</i>	<i>96.1</i>		<i>11,377,627</i>
0.99	44.998	0.357	0.845	1.879	93.0	93.0	290,332	18,479,751
		0.266	0.700	1.556	96.1	96.1	1,177,202	37,675,497
		<i>0.317</i>	<i>0.796</i>	<i>1.769</i>	<i>95.5</i>	<i>95.5</i>		<i>37,836,946</i>
0.995	51.930	0.417	0.974	1.877	93.6	93.6	441,517	28,290,323
		0.312	0.808	1.558	95.8	95.8	1,796,989	57,508,525
		<i>0.368</i>	<i>0.904</i>	<i>1.742</i>	<i>95.1</i>	<i>95.1</i>		<i>64,419,786</i>

5 CONCLUSIONS

In this paper, we presented SQSTS, a fully automated sequential procedure for providing CI estimators for steady-state quantiles of a simulation output process. SQSTS is based on the linear combination of variance-parameter estimators computed from STS and nonoverlapping batch quantiles. Initial experimentation based on the process generated by successive customer delays prior to service in a heavily initialized M/M/1 system showed that SQSTS substantially outperformed Sequest and Sequem with regard to average sample size and performed comparatively well with regard to average absolute bias, average half-length, and estimated CI coverage probability. Future work includes: (i) use of alternative weight functions for computing STS area estimators; (ii) potential enhancements for estimation of extreme quantiles ($p \notin [0.05, 0.95]$); (iii) simultaneous estimation of multiple quantiles; and (iv) extensive Monte Carlo experimentation with a variety of stochastic processes with characteristics that could pose additional challenges, e.g., multimodal marginal p.d.f.'s (Alexopoulos et al. 2018).

REFERENCES

- Alexopoulos, C., J. H. Boone, D. Goldsman, A. Lolos, K. D. Dineç, and J. R. Wilson. 2020. "Steady-State Quantile Estimation Using Standardized Time Series". In *Proceedings of the 2020 Winter Simulation Conference*, edited by K.-H. Bae, B. Feng, S. Kim, L. Lazarova-Molnar, Z. Zheng, T. Roeder, and R. Thiesing, 289–300. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.
- Alexopoulos, C., K. D. Dineç, , D. Goldsman, A. Lolos, and J. R. Wilson. 2022. "Steady-State Quantile Estimation Using Standardized Time Series". Technical report, Gebze Technical University, Georgia Institute of Technology, and North Carolina State University. <https://people.engr.ncsu.edu/jwilson/files/qestr1.pdf>, accessed 29th February 2022.
- Alexopoulos, C., D. Goldsman, A. C. Mokashi, K.-W. Tien, and J. R. Wilson. 2019. "Sequest: A Sequential Procedure for Estimating Quantiles in Steady-State Simulations". *Operations Research* 67 (4): 1162–1183. <https://people.engr.ncsu.edu/jwilson/files/sequest19or.pdf>, accessed 7th September 2019.
- Alexopoulos, C., D. Goldsman, A. C. Mokashi, and J. R. Wilson. 2017. "Automated Estimation of Extreme Steady-State Quantiles via the Maximum Transformation". *ACM Transactions on Modeling and Computer Simulation* 27 (4): 22:1–22:29.
- Alexopoulos, C., D. Goldsman, A. C. Mokashi, and J. R. Wilson. 2018. "Sequential Estimation of Steady-State Quantiles: Lessons Learned and Future Directions". In *Proceedings of the 2018 Winter Simulation Conference*, edited by M. Rabe, A. A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, 1814–1825. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.
- Alexopoulos, C., D. Goldsman, P. Tang, and J. R. Wilson. 2016. "SPSTS: A Sequential Procedure for Estimating the Steady-State Mean Using Standardized Time Series". *IIE Transactions* 48 (9): 864–880.
- Billingsley, P. 1999. *Convergence of Probability Measures*. 2nd ed. New York: John Wiley & Sons.
- Calvin, J. M., and M. K. Nakayama. 2013. "Confidence Intervals for Quantiles with Standardized Time Series". In *Proceedings of the 2013 Winter Simulation Conference*, edited by R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, 601–612. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.
- Chen, E. J., and W. D. Kelton. 2006. "Quantile and Tolerance-Interval Estimation in Simulation". *European Journal of Operational Research* 168:520–540.
- Chen, E. J., and W. D. Kelton. 2008. "Estimating Steady-State Distributions via Simulation-Generated Histograms". *Computers & Operations Research* 35 (4): 1003–1016.
- Damerdj, H. 1994. "Strong Consistency of the Variance Estimator in Steady-State Simulation Output Analysis". *Mathematics of Operations Research* 19:494–512.
- Dineç, K. D., C. Alexopoulos, D. Goldsman, A. Lolos, and J. R. Wilson. 2022a. "The Geometric-Moment Contraction Condition and the Quantile-Indicator Process". Technical report, Gebze Technical University. <https://people.engr.ncsu.edu/jwilson/files/gmc-sig-tr0722.pdf>, accessed 8th July 2022.
- Dineç, K. D., C. Alexopoulos, D. Goldsman, A. Lolos, and J. R. Wilson. 2022b. "Geometric Moment-Contraction of G/G/1 Waiting Times". In *Advances in Modeling and Simulation: Festschrift for Pierre L'Ecuyer*, edited by Z. Botev, A. Keller, C. Lemieux, and B. Tuffin. Springer. To appear. <https://people.engr.ncsu.edu/jwilson/files/gmc-gg1-tr0722.pdf>, accessed 6th July 2022.
- Dong, J., and P. W. Glynn. 2019. "The Asymptotic Validity of Sequential Stopping Rules for Confidence Interval Construction Using Standardized Time Series". In *Proceedings of the 2019 Winter Simulation Conference*, edited by N. Mustafee, K.-H. G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, 332–343. Piscataway, New Jersey: Institute of Electrical and Electronic Engineers.
- Fishman, G. S. 1978. *Principles of Discrete Event Simulation*. New York: John Wiley & Sons.
- Foley, R. D., and D. Goldsman. 1999. "Confidence Intervals Using Orthonormally Weighted Standardized Time Series". *ACM Transactions on Modeling and Simulation* 9:297–325.
- Glasserman, P. 2004. *Monte Carlo Methods in Financial Engineering*. New York: Springer-Verlag.
- Glynn, P. W., and D. L. Iglehart. 1990. "Simulation Output Analysis Using Standardized Time Series". *Mathematics of Operations Research* 15:1–16.

- Goldsman, D., M. S. Meketon, and L. W. Schruben. 1990. "Properties of Standardized Time Series Weighted Area Variance Estimators". *Management Science* 36:602–612.
- Heidelberger, P., and P. A. W. Lewis. 1984. "Quantile Estimation in Dependent Sequences". *Operations Research* 32 (1): 185–209.
- Hopp, W. J., and M. L. Spearman. 2008. *Factory Physics*. 3rd ed. New York: McGraw-Hill/Irwin.
- Meketon, M. S., and B. W. Schmeiser. 1984. "Overlapping Batch Means: Something for Nothing?". In *Proceedings of the 1984 Winter Simulation Conference*, edited by S. Sheppard, U. W. Pooch, and C. D. Pegden, 227–230. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.
- Raatikainen, K. E. E. 1990. "Sequential Procedure for Simultaneous Estimation of Several Percentiles". *Transactions of the Society for Computer Simulation* 7 (1): 21–44.
- Schruben, L. W. 1983. "Confidence Interval Estimation Using Standardized Time Series". *Operations Research* 31:1090–1108.
- Shao, X., and W. B. Wu. 2007. "Asymptotic Spectral Theory for Nonlinear Time Series". *The Annals of Statistics* 35 (4): 1773–1801.
- Shapiro, S. S., and M. B. Wilk. 1965. "An Analysis of Variance Test for Normality". *Biometrika* 52:591–611.
- von Neumann, J. 1941. "Distribution of the Ratio of the Mean Square Successive Difference to the Variance". *Annals of Mathematical Statistics* 12 (4): 367–395.
- Whitt, W. 2002. *Stochastic-Process Limits: An Introduction to Stochastic-Process Limits and Their Application to Queues*. New York: Springer.
- Wu, W. B. 2005. "On the Bahadur Representation of Sample Quantiles for Dependent Sequences". *The Annals of Statistics* 33 (4): 1934–1963.

AUTHOR BIOGRAPHIES

ATHANASIOS LOLOS is a PhD candidate in the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology. He received his BSc degree from the School of Naval Architecture and Marine Engineering at the National Technical University of Athens, Greece with highest honors. He is a recipient of a scholarship by the Alexander S. Onassis Foundation. His email address is thnlolos@gatech.edu.

J. HADEN BOONE is a PhD student in the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology. His email address is jboone31@gatech.edu.

CHRISTOS ALEXOPOULOS is a Professor in the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology. His research interests are in the areas of simulation, statistics, and optimization of stochastic systems. He is a member of INFORMS and an active participant in the Winter Simulation Conference, having been *Proceedings* Co-Editor in 1995, and Associate Program Chair in 2006. He served on the Board of Directors of WSC between 2008 and 2016. He is also an Associate Editor of *ACM Transactions on Modeling and Computer Simulation*. His e-mail address is christos@gatech.edu, and his Web page is www.isye.gatech.edu/~christos.

DAVID GOLDSMAN is a Professor in the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology. His research interests include simulation output analysis, ranking and selection, and healthcare simulation. He was Program Chair of the Winter Simulation Conference in 1995 and a member of the WSC Board of Directors between 2001–2009. His e-mail address is sman@gatech.edu, and his Web page is www.isye.gatech.edu/~sman.

KEMAL DİNÇER DİNĞEÇ is an Assistant Professor in the Department of Industrial Engineering at Gebze Technical University in Istanbul, Turkey. Previously, he was a post-doctoral researcher in the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology and at Boğaziçi University's Industrial Engineering Department. His research interests include stochastic models in manufacturing and finance, Monte Carlo simulation, and variance reduction methods for financial simulations. His email address is kdinger@yahoo.com.

ANUP C. MOKASHI is a Senior Operations Research Engineer at Memorial Sloan Kettering Cancer Center. He holds an MS in Industrial Engineering from North Carolina State University. His research interests include design and implementation of algorithms related to statistical aspects of discrete-event simulation. His career interests include applying simulation and other Operations Research techniques to large scale industrial problems. He is a member of IISE and INFORMS. His e-mail address is MokashiA@mskcc.org.

JAMES R. WILSON is a Professor Emeritus in the Edward P. Fitts Department of Industrial and Systems Engineering at North Carolina State University. His current research interests are focused on probabilistic and statistical issues in the design and analysis of simulation experiments. He is a member of ACM and ASA, and he is a Fellow of IISE and INFORMS. His email address is jwilson@ncsu.edu, and his Web page is www.ise.ncsu.edu/jwilson.