Proceedings of the 2025 Winter Simulation Conference
E. Azar, A. Djanatliev, A. Harper, C. Kogler, V. Ramamohan, A. Anagnostou, and S. J. E. Taylor, eds.

A FIXED-SAMPLE-SIZE PROCEDURE FOR ESTIMATING STEADY-STATE QUANTILES BASED ON INDEPENDENT REPLICATIONS

Athanasios Lolos[1], Christos Alexopoulos[2], David Goldsman[2], Kemal Dinçer Dingeç[3], Anup C. Mokashi[4], and James R. Wilson[5]

[1]Business Intelligence and Data Science Team, Membership and Deposit Products Division, Navy Federal Credit Union, Vienna, VA, 22180, USA
[2]H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0205, USA
[3]Department of Industrial Engineering, Gebze Technical University, 41400 Gebze, Kocaeli, TURKEY
[4]Memorial Sloan Kettering Cancer Center, 1275 York Avenue, New York NY 10065, USA
[5]Edward P. Fitts Department of Industrial and Systems Engineering, North Carolina State University, Raleigh, NC 27695-7906, USA

ABSTRACT

We introduce the first fully automated fixed-sample-size procedure (FIRQUEST) for computing confidence intervals (CIs) for steady-state quantiles based on independent replications. The user provides a dataset from a number of independent replications of arbitrary size and specifies the required quantile and nominal coverage probability of the anticipated CI. The proposed method is based on the simulation analysis methods of batching, standardized time series (STS), and sectioning. Preliminary experimentation with the waiting-time process in an M/M/1 queueing system showed that FIRQUEST performed well by appropriately handling initialization effects and delivering CIs with estimated coverage probability close to the nominal level.

1 INTRODUCTION

Steady-state simulations are an important tool for the design and performance evaluation of complex production and service systems (Law 2015). Quantiles of the marginal steady-state distribution of an output process are standard measures of risk (Glasserman 2004). Steady-state quantile estimation is a challenging problem due to (i) the potential presence of initialization bias; (ii) substantial serial correlation in the simulation output process; (iii) departures from normality; (iv) inherent bias of point estimators; and (v) the challenging nature of the marginal distribution such as nonexistence of a probability density function (p.d.f.) including discontinuities and multimodalities with sharp peaks (Alexopoulos et al. 2018).

In their recent work, Alexopoulos et al. (2019) and Lolos et al. (2022, 2024) introduced Sequest and SQSTS, respectively, two automated sequential methods for steady-state quantile estimation. However, users often (i) use simulation models not integrated with the underlying sequential method or (ii) encounter datasets that are limited due to budget constraints.

To address these issues, Lolos et al. (2023, 2025) developed FQUEST, an automated fixed-sample-size procedure for steady-state quantile estimation based on a single run. FQUEST stands out due to its following unique characteristics in which the procedure: (i) utilizes the STS methodology; (ii) handles effectively the simulation initialization problem; and (iii) provides a warning when the provided dataset is insufficient and, subject to the user's approval, delivers a heuristic CI.

Steady-state analysis methods based on a single simulation replication are convenient in the sense that only data from the onset of the run may have to be eliminated to diminish the effects of initialization bias. Unfortunately, the potential of pronounced autocorrelation in the underlying output process may require

979-8-3315-8726-0/25/$31.00 ©2025 IEEE

excessively large sample sizes to attenuate this correlation effect and yield reliable CIs for the performance measure of interest. On the other hand, steady-state estimation methods based on independent replications are convenient and reduce the correlation problems. For practical purposes, the need for replications is further enhanced by the fact that the replications can be executed simultaneously on different cores/threads within a single computer or on different computers on a network, provided that the software being used for simulation supports this (Law 2015). On the negative side, independent replications can induce systematic bias if insufficient truncation is applied at the onset of each replication (Alexopoulos and Goldsman 2004, Fishman 2001). Further, for fixed-sample-size procedures, one has to decide on the number of replications and the run length within each replication.

In this paper, we introduce FIRQUEST, the first fully automated, fixed-sample-size method for estimating steady-state quantiles based on independent replications. FIRQUEST is essentially an extension of the FQUEST procedure of Lolos et al. (2025) with adjustments to handle replicated sample paths and more aggressive steps for removing potential warm-up effects that can induce a systematic bias across replicate estimates (Alexopoulos and Goldsman 2004).

FIRQUEST provides a CI and a point estimate for a selected steady-state quantile, with a user-specified error probability. The user feeds FIRQUEST with datasets of the same size produced by an arbitrary number of independent replications. The theory on which the CIs used in FIRQUEST are based can be found in Alexopoulos et al. (2020, 2025) and in Lolos et al. (2024).

In Section 2, we outline the notation, main assumptions, and theoretical background of FIRQUEST. In Section 3, we present the algorithm. In Section 4, we conduct a preliminary performance evaluation of FIRQUEST based on the waiting-time process in an M/M/1 system. Section 5 contains a summary of our work and discussions about future research directions.

2 FOUNDATIONS

In this section, we outline the notation, assumptions, and core results that form the basis of FIRQUEST.

2.1 Notation

The set of real numbers is denoted by $\mathbb{R} \equiv (-\infty, \infty)$ while the set of integers is denoted by $\mathbb{Z} \equiv \{0, \pm 1, \pm 2, \ldots\}$. For $p \in (0,1)$, the p-quantile of a random variable (r.v.) Y with cumulative distribution function (c.d.f.) $F(y)$ is defined as $y_p \equiv F^{-1}(p) \equiv \inf\{y : F(y) \geq p\}$. Let $\{Y_k : k \geq 0\}$ be a discrete-time stationary simulation output process with marginal c.d.f. $F(\cdot)$. Our goal is to compute a point estimate and a CI for the quantile y_p based on a finite sample $\{Y_k : k = 1, \ldots, n\}$ that is free of initialization bias. Let $Y_{(1)} \leq \cdots \leq Y_{(n)}$ be the respective order statistics. The classical point estimator of y_p is $\widetilde{y}_p(n) \equiv Y_{(\lceil np \rceil)}$ (the empirical p-quantile), where $\lceil \cdot \rceil$ denotes the ceiling function.

For all $k \geq 1$ and $y \in \mathbb{R}$, we define the indicator r.v. $I_k(y) \equiv 1$ if $Y_k \leq y$, and $I_k(y) \equiv 0$ otherwise; hence $\mathrm{E}[I_k(y_p)] = p$. We let $\overline{I}_n(y) \equiv n^{-1} \sum_{k=1}^{n} I_k(y)$, for $n \geq 1$; and for each $\ell \in \mathbb{Z}$, we let $\rho_{I(y)}(\ell) \equiv \mathrm{Corr}[I_k(y), I_{k+|\ell|}(y)]$ denote the autocorrelation function of the indicator process $\{I_k(y) : k \geq 0\}$ at lag ℓ. Throughout the paper we use the following notation as well: Z denotes an r.v. having the standard normal distribution, $N(0,1)$; for each integer $\nu \geq 1$, $\boldsymbol{Z}_\nu \equiv \left[Z_1, \ldots, Z_\nu\right]^{\mathsf{T}}$ denotes a $\nu \times 1$ vector whose components are independent and identically distributed (i.i.d.) $N(0,1)$ r.v.'s; χ^2_ν denotes a chi-squared r.v. with ν degrees of freedom (d.f.); t_ν denotes an r.v. having Student's t-distribution with ν d.f.; and for each $\delta \in (0,1)$, $t_{\delta,\nu}$ denotes the δ-quantile of t_ν.

The basic (unadjusted) $100(1-\alpha)\%$ CIs for y_p have the form

$$\widetilde{y}_p(n) \pm t_{1-\alpha/2,\nu}\, \widehat{\sigma}_{\widetilde{y}_p} \big/ \sqrt{n},$$

where $\widehat{\sigma}^2_{\widetilde{y}_p}$ is an estimator of the (quantile) variance parameter $\sigma^2_{\widetilde{y}_p} \equiv \lim_{n\to\infty} n\mathrm{Var}\big[\widetilde{y}_p(n)\big]$, the d.f. ν is determined by the underlying quantile-estimation method, and $\alpha \in (0,1)$ is selected by the user (Asmussen and Glynn 2007).

2.2 Assumptions

This subsection contains the main assumptions for the processes $\{Y_k : k \geq 0\}$ and $\{I_k(y_p) : k \geq 0\}$.

Geometric-Moment Contraction (GMC) Condition (Wu 2005). The process $\{Y_k : k \geq 0\}$ is defined by a function $\xi(\cdot)$ of a sequence of i.i.d. r.v.'s $\{\varepsilon_j : j \in \mathbb{Z}\}$ such that $Y_k = \xi(\ldots, \varepsilon_{k-1}, \varepsilon_k)$ for $k \geq 0$. Moreover, there exist constants $\psi > 0$, $C_\psi > 0$, and $r_\psi \in (0,1)$ 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

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

The GMC condition holds for a wide range of random processes (Shao and Wu 2007; Alexopoulos et al. 2019). Alexopoulos et al. (2025) and Dingeç et al. (2025) elaborate on the applicability of the GMC condition compared to several alternatives. Recently, Dingeç et al. (2022) showed that the customer waiting-time processes in an M/M/1 queueing system and a G/G/1 system with non-heavy-tailed service-time distributions satisfy the GMC condition.

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, at the quantile of interest y_p, we have $f(y_p) > 0$, and the derivative $f'(y_p)$ exists.

Short-Range Dependence (SRD) of the Indicator Process. The indicator process $\{I_k(y) : k \geq 0\}$ has the SRD property so that for all $y \in \mathbb{R}$,

$$0 \leq \sum_{\ell \in \mathbb{Z}} \rho_{I(y)}(\ell) \leq \sum_{\ell \in \mathbb{Z}} \left|\rho_{I(y)}(\ell)\right| < \infty.$$

Thus the variance parameters associated with the random processes $\{\overline{I}_n(y_p)\}$ and $\{\widetilde{y}_p(n)\}$, respectively, satisfy the relations

$$\sigma^2_{I(y_p)} \equiv \lim_{n\to\infty} n\mathrm{Var}\big[\overline{I}_n(y_p)\big] = p(1-p)\sum_{\ell \in \mathbb{Z}} \rho_{I(y_p)}(\ell) \in [0, \infty),$$
$$\sigma^2_{\widetilde{y}_p} = \lim_{n\to\infty} n\mathrm{Var}\big[\widetilde{y}_p(n)\big] = \frac{\sigma^2_{I(y_p)}}{f^2(y_p)} \in [0, \infty).$$

Functional Central Limit Theorem (FCLT) for the Indicator Process. Let $D \equiv D[0,1]$ denote the space of real-valued functions on $[0,1]$ that are right continuous with left-hand limits; let $\mathcal{W}$ denote a standard Brownian motion on $[0,1]$; and let $\underset{n\to\infty}{\Longrightarrow}$ denote weak convergence as $n \to \infty$ (Billingsley 1999). Also, let $\lfloor\cdot\rfloor$ denote the floor function. We define the following sequence of random functions $\{\mathcal{I}_n : n \geq 1\}$ in D,

$$\mathcal{I}_n(t) \equiv \mathcal{I}_n(t; y_p) \equiv \frac{\lfloor nt \rfloor}{\sigma_{I(y_p)} n^{1/2}}\big[\overline{I}_{\lfloor nt \rfloor}(y_p) - p\big], \quad \text{for } t \in [0,1] \text{ and } n \geq 1,$$

We assume that this random-function sequence satisfies the FCLT, i.e., $\mathcal{I}_n \underset{n\to\infty}{\Longrightarrow} \mathcal{W}$.

The literature contains sound theoretical and empirical evidence of the mutual compatibility of the GMC, DR, SRD, and FCLT conditions (Whitt 2002; Dingeç et al. 2025).

2.3 Asymptotic Properties Based on Nonoverlapping Batches

FIRQUEST is based on nonoverlapping batches. Given a fixed batch count $b \geq 2$, for $j = 1, \ldots, b$ the jth nonoverlapping batch of size $m \geq 1$ consists of the subsequence $\{Y_{(j-1)m+1}, \ldots, Y_{jm}\}$, where we assume $n = bm$. The batch mean of the indicator r.v.'s associated with the jth batch is $\overline{I}_{j,m}(y) \equiv m^{-1}\sum_{\ell=1}^m I_{(j-1)m+\ell}(y)$.

Similarly to the full-sample case, we define the order statistics $Y_{j,(1)} \leq \cdots \leq Y_{j,(m)}$ corresponding to the jth batch. Then the jth batched quantile estimator (BQE) of y_p is $\widehat{y}_p(j,m) \equiv Y_{j,(\lceil mp \rceil)}$. Alexopoulos et al. (2019) establish that

$$m^{1/2}\big[\widehat{y}_p(1,m)-y_p,\ldots,\widehat{y}_p(b,m)-y_p\big]^{\mathsf{T}} \underset{m\to\infty}{\Longrightarrow} \sigma_{\widetilde{y}_p}\mathbf{Z}_b.$$

2.4 Confidence Intervals for Quantiles

FIRQUEST utilizes CIs computed from STSs based on nonoverlapping batches, the BQEs $\widehat{y}_p(j,m)$, and the full-sample empirical quantile $\widetilde{y}_p(n)$. We define the full-sample STS process for quantile estimation as

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

where $\widetilde{y}_p(\lfloor nt \rfloor)$ is the empirical p-quantile (i.e., the $\lceil p\lfloor nt \rfloor \rceil$-th order statistic) computed from the partial sample $\{Y_k : k=1,\ldots,\lfloor nt \rfloor\}$. Further, Alexopoulos et al. (2025) show that if $\{Y_k : k \geq 1\}$ satisfies the assumptions of Theorem 1 therein, then

$$\big[n^{1/2}\big(\widetilde{y}_p(n)-y_p\big),T_n\big] \underset{n\to\infty}{\Longrightarrow} \sigma_{\widetilde{y}_p}\big[\mathcal{W}(1),\mathcal{B}\big],$$

where $\mathcal{B}(t) \equiv \mathcal{W}(t)-t\mathcal{W}(1)$ for $t \in [0,1]$ is a standard Brownian bridge process that is independent of $\mathcal{W}(1)$. Hence, the full-sample r.v. $n^{1/2}\big(\widetilde{y}_p(n)-y_p\big)$ is asymptotically independent of T_n as $n \to \infty$.

The full-sample STS area estimator of the variance parameter $\sigma^2_{\widetilde{y}_p}$ is $A_p^2(w;n)$, where:

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

and $w(\cdot)$ is a deterministic weight function that is bounded and continuous a.e. on $[0,1]$; and the r.v.

$$Z(w) \equiv \int_0^1 w(t)\mathcal{B}(t)\,\mathrm{d}t \sim N(0,1). \tag{1}$$

The r.v. $Z(w)$ is the signed, weighted area bounded by the random function $w(t)\mathcal{B}(t)$ for $t \in [0,1]$ and the horizontal axis so that $Z(w)$ is normally distributed. Moreover, Alexopoulos et al. (2025) show that the STS area estimators $\{A_p(w;n) : n \geq 1\}$ for $\sigma^2_{\widetilde{y}_p}$ satisfy the following weak-convergence results:

$$A_p(w;n) \underset{n\to\infty}{\Longrightarrow} \sigma_{\widetilde{y}_p}Z(w) \quad \text{and} \quad A_p^2(w;n) \underset{n\to\infty}{\Longrightarrow} \sigma^2_{\widetilde{y}_p}\chi^2_1.$$

Several weight functions satisfy condition (1) including the constant $w_0(t) \equiv \sqrt{12}$ (Schruben 1983), the quadratic $w_2(t) \equiv \sqrt{840}(3t^2-3t+1/2)$ (Goldsman et al. 1990), and the orthonormal family $\{w_{\cos,j}(t) \equiv \sqrt{8}\pi j\cos(2\pi jt) : j=1,2,\ldots\}$ (Foley and Goldsman 1999). Since none of the latter weight functions have provably yielded less-biased estimators (based on STS) than $w_0(\cdot)$ (Lolos et al. 2024), we will use the constant weight $w_0(\cdot)$ for the performance evaluation.

Below we describe the extension of the aforementioned results for the case of nonoverlapping batches of size m. For $j=1,\ldots,b$, we define $\widehat{y}_p(j,\lfloor mt \rfloor)$ as the empirical p-quantile (i.e., the $\lceil p\lfloor mt \rfloor \rceil$-th order statistic) computed from the partial sample $\{Y_{(j-1)m+k} : k=1,\ldots,\lfloor mt \rfloor\}$, and the STS-based quantile-estimation process formed from batch j as

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

Moreover, we define the signed area computed from batch j as

$$A_p(w;j,m) \equiv m^{-1} \sum_{k=1}^{m} w(k/m) T_{j,m}(k/m).$$

Lolos et al. (2023, 2025) provide all the established theoretical results that form the basis of the sequential and fixed-sample-size procedures therein. These results are also utilized by FIRQUEST, with the main difference that an equal number of batches is formed within each replicate path.

For simplicity, assume that we have generated R i.i.d. stationary sample paths of the process $\{Y_k : k \geq 1\}$, each of size bm, so that $N = Rbm$. We split each replicate path in b nonoverlapping batches of size m each. From each batch we compute the respective empirical quantile and weighted signed area. Also, we denote the replicate batched quantile estimator (RBQE) as $\{\widehat{y}_p(j,m) : j = 1,\ldots,Rb\}$ and the (replicate) signed areas as $\{A_p(w;j,m) : j = 1,\ldots,Rb\}$, where the subscript j in $\widehat{y}_p(j,m)$ or $A_p(w;j,m)$ denotes the ith RBQE or signed area, respectively, from replication $r+1$ with $r = \lfloor j/b \rfloor$ and $i \equiv j - rb$. For example, if $b = 20$, $\widehat{y}_p(43,m)$ is the 3rd RBQE from replication 3. Also, let $\widetilde{y}_p(N)$ be the empirical quantile computed from the entire dataset comprised of the R sample paths.

We define the replicated batched STS area estimator as

$$\mathcal{A}_p(w;R,b,m) \equiv (Rb)^{-1} \sum_{j=1}^{Rb} A_p^2(w;j,m).$$

We also define the average RBQE

$$\overline{\overline{y}}_p(R,b,m) \equiv (Rb)^{-1} \sum_{j=1}^{Rb} \widehat{y}_p(j,m)$$

and the "average" squared deviations of the RBQEs away from the average RBQE $\overline{\overline{y}}_p(R,b,m)$ and the full-sample quantile estimator $\widetilde{y}_p(N)$, respectively,

$$S_p^2(R,b,m) \equiv (Rb-1)^{-1} \sum_{j=1}^{Rb} \left[\widehat{y}_p(j,m) - \overline{\overline{y}}_p(R,b,m)\right]^2 \quad \text{and}$$

$$\widetilde{S}_p^2(R,b,m) \equiv (Rb-1)^{-1} \sum_{j=1}^{Rb} \left[\widehat{y}_p(j,m) - \widetilde{y}_p(N)\right]^2.$$

Finally, we let

$$\mathcal{N}_p(R,b,m) \equiv m S_p^2(R,b,m) \quad \text{and} \quad \widetilde{\mathcal{N}}_p(R,b,m) \equiv m \widetilde{S}_p^2(R,b,m),$$

and we define the combined estimators of the variance parameter $\sigma_{\widetilde{y}_p}^2$ associated with the quantile process

$$\mathcal{V}_p(w;R,b,m) \equiv \frac{Rb\mathcal{A}_p(w;R,b,m) + (Rb-1)\mathcal{N}_p(R,b,m)}{2Rb-1} \quad \text{and}$$

$$\widetilde{\mathcal{V}}_p(w;R,b,m) \equiv \frac{Rb\mathcal{A}_p(w;R,b,m) + (Rb-1)\widetilde{\mathcal{N}}_p(R,b,m)}{2Rb-1}.$$

Based on the results of Alexopoulos et al. (2025) and Lolos et al. (2025), we can easily show that each of the (Rb)-dimensional random vectors $\left[\widehat{y}_p(1,m),\ldots,\widehat{y}_p(Rb,m)\right]^\mathsf{T}$ and $\left[A_p(w;1,m),\ldots,A_p(w;Rb,m)\right]^\mathsf{T}$ converges to a vector of i.i.d. normal r.v.'s as $m \to \infty$. Hence, one can readily see that, for fixed R and b,

$$\mathcal{A}_p(w;R,b,m) \underset{m\to\infty}{\Longrightarrow} \sigma_p^2 \chi_{Rb}^2/(Rb). \tag{2}$$

Suppose that R and b are fixed. The asymptotic validity, as $m \to \infty$, of the $100(1-\alpha)\%$ CI in Equation (3) follows immediately from Equation (2) above. We postulate that the remaining CIs in Equations (3)–(7) below are also asymptotically valid as $m \to \infty$:

$$\widetilde{y}_p(N) \pm t_{1-\alpha/2,Rb}\big[\mathcal{A}_p(w;R,b,m)/N\big]^{1/2}, \tag{3}$$

$$\overline{\overline{y}}_p(R,b,m) \pm t_{1-\alpha/2,Rb}\big[\mathcal{A}_p(w;R,b,m)/N\big]^{1/2}, \tag{4}$$

$$\widetilde{y}_p(N) \pm t_{1-\alpha/2,Rb-1}\big[\widetilde{\mathcal{N}}_p(R,b,m)/N\big]^{1/2}, \tag{5}$$

$$\overline{\overline{y}}_p(R,b,m) \pm t_{1-\alpha/2,Rb-1}\big[\widetilde{\mathcal{N}}_p(R,b,m)/N\big]^{1/2}, \tag{6}$$

and

$$\widetilde{y}_p(N) \pm t_{1-\alpha/2,2Rb-1}\big[\widetilde{\mathcal{V}}_p(w;R,b,m)/N\big]^{1/2}. \tag{7}$$

If the dataset is deemed to be appropriate, FIRQUEST will deliver the CI in Equation (7) above. Otherwise, subject to the user's approval, the method also constructs approximate CIs from the RBQEs $\{\widehat{y}_p(j,m) : j = 1,\ldots,Rb\}$ and the full-sample estimator $\widetilde{y}_p(N)$ using an adjustment for residual skewness in the RBQEs. Since the RBQEs are not computed from a single sample path, we do not employ the additional adjustment for residual autocorrelation in Alexopoulos et al. (2019) and Lolos et al. (2023).

Initially, we calculate the sample skewness of the RBQEs

$$\widehat{B}_{\widehat{y}_p}(R,b,m) \equiv \frac{Rb}{(Rb-1)(Rb-2)}\sum_{j=1}^{Rb}\left[\frac{\widehat{y}_p(j,m)-\overline{\overline{y}}_p(R,b,m)}{S_p(R,b,m)}\right]^3,$$

we compute the skewness-adjustment parameter

$$\vartheta \equiv \frac{\widehat{B}_{\widehat{y}_p}(R,b,m)}{6\sqrt{Rb}},$$

and define the skewness-adjustment function

$$G(\zeta) \equiv \begin{cases} \zeta & \text{if } |\vartheta| \le 0.001, \\ \frac{[1+6\vartheta(\zeta-\vartheta)]^{1/3}-1}{2\vartheta} & \text{if } |\vartheta| > 0.001, \end{cases}$$

for all real ζ. Then we set

$$G_1 \equiv G(t_{1-\alpha/2,Rb-1})\sqrt{\widetilde{S}_p^2(R,b,m)/(Rb)}, \quad \text{and} \quad G_2 \equiv G(t_{\alpha/2,Rb-1})\sqrt{\widetilde{S}_p^2(R,b,m)/(Rb)}.$$

The (asymmetric) skewness-adjusted CI for y_p is given by

$$\left[\min\left(\widetilde{y}_p(N)-G_1,\widetilde{y}_p(N)-G_2\right),\max\left(\widetilde{y}_p(N)-G_1,\widetilde{y}_p(N)-G_2\right)\right]. \tag{8}$$

3 THE FIRQUEST ALGORITHM

In this section we present the FIRQUEST procedure. FIRQUEST uses the same batching scheme in each replicate path based on b nonoverlapping batches of size m each, with these parameters updated based on outcomes of statistical tests for independence and normality. At a high level, FIRQUEST is comprised of four main blocks. Below we provide a synopsis of each block.

The first block consists of Steps **[0]**–**[2]** which initialize the experimental parameters. The user provides the number of independent replications R, the fixed size n of each replication, the probability p corresponding to the quantile of interest, and the nominal error probability $\alpha \in (0,1)$ for the CI for y_p. We also define an array, denoted by $\boldsymbol{s}$, of batch counts for Steps **[5]**–**[9]** as a function of the number of independent replications R, and we set q equal to the number of elements in $\boldsymbol{s}$. The assignment of the elements of $\boldsymbol{s}$ is based on the following guidelines: (i) start with a total batch count $Rb \geq 16$ and keep $Rb \geq 10$; (ii) use the same number of batches from every replication; (iii) use at least one batch from every replication; and (iv) if $R < 33$, use $Rb \leq 66$ total batches.

The second block includes Steps **[3]**–**[5]** and deals with the potential transient effects in each replication. Step **[3]** consists of a loop that tests the signed areas $\{A_p(w;j,m)\}$ in each replication computed from the first bm observations for independence (the tail $n - bm$ observations are ignored, but not discarded) using a two-sided test based on von Neumann's ratio (von Neumann 1941) with progressively decreasing significance level $\beta\psi(\ell)$ on iteration ℓ. Every time the randomness test fails, we keep increasing the batch size $m = \lfloor n/b \rfloor$ until the updated sample size exceeds n. If the randomness test fails with the largest allowable batch size $\lfloor n/b \rfloor$ even within one of the independent replications, FIRQUEST exits Step **[3]** and moves to Step **[4]**, where it issues a warning to the user regarding the insufficiency of the length of a replication and seeks permission to construct a heuristic conservative CI. In Step **[5]**, FIRQUEST calculates $m_{\max}$, the maximum batch size m that was used across the independent replications in Step **[3]**. Then, it removes the $m_{\max}$ first observations from every replication, sets the new run length to $n^* = n - m_{\max}$, and reindexes the truncated dataset in each replication.

The third block consists of Steps **[6]**–**[9]**. Here, we conduct the von Neumann (1941) test again and the one-sided test of Shapiro and Wilk (1965) for univariate normality to assess the convergence of each of the replicate signed areas $\{A_p(w;j,m) : j = 1,\ldots,Rb\}$ and the RBQEs $\{\widehat{y}_p(j,m) : j = 1,\ldots,Rb\}$ to asymptotic independence and normality. Each of the Steps **[6]**–**[9]** has a very similar structure. First we compute the replicate signed areas $\{A_p(w;j,m) : j = 1,\ldots,Rb\}$ or the RBQEs $\{\widehat{y}_p(j,m) : j = 1,\ldots,Rb\}$ and conduct the pertinent statistical test using the fixed significance level of $\beta = 0.3$. This relatively large significance level is kept constant to avoid passing a test with an inadequately small batch size leading to unreliable CIs. If the pertinent test is passed, FIRQUEST proceeds to the next step; otherwise, the batch count in each replication decreases to the next element of the array $\boldsymbol{s}$. Since q is equal to the number of elements in $\boldsymbol{s}$, we can have up to q failed attempts to pass any of the statistical tests in Steps **[6]**–**[9]**. If at any point a statistical test fails with $v = q$, then FIRQUEST skips the remaining statistical tests and moves to Step **[10]**.

Finally, the last block consists of Step **[10]**: If the statistical tests within the third block are passed, the procedure delivers the CI in Equation (7) based on the combined variance estimator. Otherwise, it delivers a conservative CI, subject to the user's approval. This heuristic CI is wider by construction to meet the user's coverage probability requirements even in unfavorable circumstances with inadequate provided replication sizes.

The formal algorithmic statement of FIRQUEST follows. We present the algorithm for a general weight function $w(\cdot)$ satisfying Equation (1), but we use $w_0(t) = \sqrt{12}$ for the results in Section 4.

Algorithm FIRQUEST

---

[0] User-Initialization: A number of R replicate sample paths, each of size n (total sample size Rn), the probability p corresponding to the desired quantile, and the error probability $\alpha \in (0,1)$.

[1] Parameter-Initialization: Set the number of batches $b = 25$, the batch size $m = 500$, $\ell = 1$, $v = 1$, and `flag` = `false`. Also set $\beta = 0.30$ and

$$\boldsymbol{s} = \begin{cases} [14,11,8,5] & \text{if } R = 2, \\ [10,8,6,4] & \text{if } R = 3, \\ [6,5,4,3] & \text{if } R = 4, \\ [5,4,3,2] & \text{if } 5 \le R < 10, \\ [4,3,2,1] & \text{if } 10 \le R < 17, \\ [3,2,1] & \text{if } 17 \le R < 23, \\ [2,1] & \text{if } 23 \le R < 33, \\ [1] & \text{if } 33 \le R. \end{cases}$$

Further, set q equal to the number of elements in $\boldsymbol{s}$. Let $w(t)$, $t \in [0,1]$, be the weight function and define the initial significance level for the first hypothesis test in Step **[3]** as $\beta\psi(\ell)$, where $\psi(\ell) \equiv \exp\left[-\eta(\ell-1)^{\theta}\right]$, $\ell = 1,2,\ldots$, with $\eta = 0.2$ and $\theta = 2.3$.

[2] **If** $n < bm$:

Set $m \leftarrow \lfloor n/b \rfloor$;

[3] For each replicate path execute the following loop:

Until von Neumann's test fails to reject randomness **or** `flag` = `true`:

- Compute the signed areas $\{A_p(w;j,m)\}$ from the current replicate path;
- Assess the randomness of the signed areas $\{A_p(w;j,m)\}$ from the current replication using von Neumann's two-sided randomness test with significance level $\beta\psi(\ell)$;
- Set $\ell \leftarrow \ell+1$ and $m \leftarrow [\![m\sqrt{2}]\!]$;
- **If** $n < bm$ **and** $m \neq \lfloor n/b \rfloor$:

  Set $m \leftarrow \lfloor n/b \rfloor$;

  Else

  Set `flag` $\leftarrow$ `true`;

Set $\ell \leftarrow 1$ and $m \leftarrow 500$.

[4] If the randomness test in Step **[3]** failed for any of the independent replications due to insufficient length, then issue a warning and seek permission from the user to continue with the construction of a CI. If the user declines, then exit without delivering a CI.

[5] Remove the first $m_{\max}$ observations from each replication, reindex the truncated datasets, and set n^* $(= n - m_{\max})$ equal to the size of the truncated sample of each replication. Set the number of batches $b \leftarrow s[v]$ and calculate the batch size as $m \leftarrow \lfloor n^*/b \rfloor$. Ignore the initial $n^* - bm$ observations from each replication.

[6] **Until** von Neumann's test fails to reject randomness **or** $v = q+1$ (a test has failed with the minimum allowable number of batches in $\boldsymbol{s}$):

- Compute the replicate signed areas $\{A_p(w;j,m) : j = 1,\ldots,Rb\}$ across all replications;
- Assess the randomness of the replicate signed areas $\{A_p(w;j,m) : j = 1,\ldots,Rb\}$ using von Neumann's two-sided randomness test with significance level β;
- Set $v \leftarrow v+1$. Update $b \leftarrow s[v]$ and $m \leftarrow \lfloor n^*/b \rfloor$. Ignore the initial $n^* - bm$ observations from each replication.

[7] **Until** the Shapiro-Wilk test fails to reject normality **or** $v = q+1$ (a test has failed with the minimum allowable number of batches in $\boldsymbol{s}$):

- Compute the replicate signed areas $\{A_p(w;j,m) : j = 1,\ldots,Rb\}$;
- Assess the univariate normality of the replicate signed areas $\{A_p(w;j,m) : j = 1,\ldots,Rb\}$ using the Shapiro–Wilk test with significance level β;

- Set $v \leftarrow v+1$. Update $b \leftarrow s[v]$ and $m \leftarrow \lfloor n^*/b \rfloor$. Ignore the initial $n^* - bm$ observations from each replication.

[8] **Until** von Neumann's test fails to reject randomness **or** $v = q+1$ (a test has failed with minimum allowable number of batches in $\boldsymbol{s}$):

- Compute the RBQEs $\{\widehat{y}_p(j,m) : j = 1,\ldots,Rb\}$;
- Assess the randomness of the RBQEs $\{\widehat{y}_p(j,m) : j = 1,\ldots,Rb\}$ using von Neumann's two-sided randomness test with significance level β;
- Set $v \leftarrow v+1$. Update $b \leftarrow s[v]$ and $m \leftarrow \lfloor n^*/b \rfloor$. Ignore the initial $n^* - bm$ observations from each replication.

[9] **Until** the Shapiro–Wilk test fails to reject normality **or** $v = q+1$ (a test has failed with minimum allowable number of batches in $\boldsymbol{s}$):

- Compute the RBQEs $\{\widehat{y}_p(j,m) : j = 1,\ldots,Rb\}$;
- Assess the univariate normality of the BQEs $\{\widehat{y}_p(j,m) : j = 1,\ldots,Rb\}$ using the Shapiro–Wilk test with significance level β;
- Set $v \leftarrow v+1$. Update $b \leftarrow s[v]$ and $m \leftarrow \lfloor n^*/b \rfloor$. Ignore the initial $n^* - bm$ observations from each replication.

[10] Set $N^* \leftarrow Rbm$.

If $v < q+1$ (no statistical test in Steps **[6]**–**[9]** failed):

- Compute the combined variance estimator $\widetilde{\mathcal{V}}_p(w;R,b,m)$, deliver the $100(1-\alpha)\%$ CI in Equation (7) for y_p, and exit;

Else

- Issue a warning that a statistical test failed due to insufficiency of the dataset and seek permission from the user to continue with the construction of a CI. If the user declines, then exit without delivering a CI;
- Compute

$$h_{\alpha,R,b,m} = \max\left\{ t_{1-\alpha/2,Rb}\sqrt{\frac{\mathcal{A}_p(w;R,b,m)}{N^*}}, t_{1-\alpha/2,Rb-1}\sqrt{\frac{\widetilde{\mathcal{N}}_p(R,b,m)}{N^*}} \right\}.$$

Then, construct the following approximate CIs for y_p with HL $h_{\alpha,R,b,m}$:

$$\widetilde{y}_p(N^*) \pm h_{\alpha,R,b,m} \quad \text{and} \quad \overline{\overline{y}}_p(R,b,m) \pm h_{\alpha,R,b,m}, \tag{9}$$

with the first CI centered around the full-sample point estimator $\widetilde{y}_p(N^*)$ and the second centered around the average RBQE $\overline{\overline{y}}_p(R,b,m) = (Rb)^{-1}\sum_{j=1}^{Rb}\widehat{y}_p(j,m)$;

- Construct the (asymmetric) skewness-adjusted CI

$$\left[\min\left(\widetilde{y}_p(N^*) - G_1, \widetilde{y}_p(N^*) - G_2\right), \max\left(\widetilde{y}_p(N^*) - G_1, \widetilde{y}_p(N^*) - G_2\right)\right] \tag{10}$$

with G_1 and G_2 defined in Equation (8);

- Deliver the full-sample point estimator $\widetilde{y}_p(N^*)$ and the smallest interval containing the CIs in Equations (9) and (10) above, and exit.

4 EXPERIMENTAL RESULTS

In this section we present an initial empirical evaluation of FIRQUEST based on the waiting-time sequence in an M/M/1 queueing system with arrival rate $\lambda = 0.9$, service rate $\omega = 1$ (traffic intensity $\rho = 0.9$), and First In, First Out (FIFO) service discipline. To assess the ability of the FIRQUEST method to deal with the transient phase, we initiate the system in the empty-and-idle state. For the experimental results we used the constant weight function $w_0(\cdot)$. Extensive numerical results, comparisons with FQUEST (using identical total sample sizes), and suggestions for selecting the number of replications R are in Lolos (2023).

Table 1: Experimental results for FIRQUEST with $R \in \{5, 10\}$ replications with regard to point and 95% CI estimation of y_p for the waiting-time process in an M/M/1 system with traffic intensity 0.9 initialized in the empty-and-idle state based on 1,000 independent replications.

p	y_p	R	Repl. Size	Point Est.	Avg. |Bias|	Avg. 95% CI HL	Avg. 95% CI rel. prec. (%)	Avg. 95% CI cov. (%)	$\overline{m}$	$\overline{b}$	St. Dev. HL	Avg. Trunc. Point
0.5	5.878	5	40,000	5.886	0.191	0.671	11.377	96.5	16,125	13.28	0.388	955
		10	20,000	5.874	0.195	0.686	11.645	94.6	16,861	12.89	0.409	787
		5	100,000	5.883	0.122	0.373	6.330	95.7	37,548	14.78	0.176	957
		10	50,000	5.880	0.126	0.375	6.371	94.8	37,382	16.27	0.186	1,094
		5	200,000	5.885	0.085	0.249	4.238	95.2	70,601	15.89	0.093	957
		10	100,000	5.877	0.085	0.244	4.145	95.7	66,135	19.34	0.100	1,093
0.7	10.986	5	40,000	10.999	0.365	1.296	11.748	95.9	16,284	13.17	0.788	969
		10	20,000	10.977	0.368	1.302	11.817	95.1	17,004	12.72	0.817	788
		5	100,000	10.990	0.227	0.720	6.546	95.9	37,751	14.74	0.364	972
		10	50,000	10.988	0.236	0.711	6.467	94.9	37,858	16.03	0.357	1,116
		5	200,000	10.996	0.162	0.479	4.355	95.4	72,044	15.52	0.186	972
		10	100,000	10.982	0.161	0.473	4.301	95.3	68,577	18.62	0.207	1,116
0.9	21.972	5	40,000	22.008	0.902	3.493	15.764	94.9	17,318	12.09	2.700	922
		10	20,000	21.940	0.888	3.465	15.675	94.7	17,708	11.85	2.642	782
		5	100,000	21.982	0.550	1.835	8.335	95.6	40,422	13.54	1.067	924
		10	50,000	21.969	0.571	1.856	8.438	94.1	41,260	13.93	1.111	1,049
		5	200,000	21.993	0.389	1.183	5.378	94.9	75,094	14.93	0.558	924
		10	100,000	21.956	0.391	1.194	5.430	94.4	76,012	16.24	0.616	1,049
0.95	28.904	5	40,000	28.964	1.380	6.028	20.616	94.2	17,958	11.51	4.972	890
		10	20,000	28.858	1.374	5.889	20.169	94.5	18,335	11.01	4.846	777
		5	100,000	28.935	0.844	3.084	10.612	95.3	42,070	12.81	2.124	891
		10	50,000	28.898	0.881	3.007	10.383	94.1	43,657	12.57	1.920	1,006
		5	200,000	28.940	0.596	1.896	6.546	95.6	79,224	13.99	0.963	891
		10	100,000	28.884	0.597	1.905	6.584	95.0	81,555	14.49	1.095	1,006
0.99	44.998	5	40,000	45.171	3.349	12.026	25.964	90.8	18,953	10.56	8.621	858
		10	20,000	44.920	3.314	11.360	24.650	90.8	19,074	10.16	8.079	773
		5	100,000	45.116	2.139	9.051	19.812	92.8	46,531	11.13	6.797	858
		10	50,000	44.950	2.138	8.900	19.577	94.2	47,671	10.63	6.387	969
		5	200,000	45.083	1.484	6.206	13.669	94.5	90,607	11.66	4.806	858
		10	100,000	44.989	1.509	6.102	13.448	94.0	93,956	11.18	4.702	969

Table 1 above displays the experimental results. For the evaluation we used three total sample sizes $N \in \{200{,}000, 500{,}000, 1{,}000{,}000\}$, two replication counts $R \in \{5, 10\}$, and the nominal 95% ($\alpha = 0.05$) CI coverage probability. All entries were computed from 1,000 independent trials. In particular, column 1 contains selected values of p, while column 2 displays the exact value of the respective quantile y_p, which was computed by inverting the c.d.f. in Equation (36) of Alexopoulos et al. (2019). Columns 3 and 4 contain the number of independent replications R and the replication size n, respectively ($n = N/R$). Column 5 displays $\widetilde{y}_p(n^*)$, the average value of the point estimate. Column 6 lists the average value of the absolute error $|\widetilde{y}_p(n^*) - y_p|$. In Columns 7–9, we provide the average value of the half-length (HL) of the 95% CI for y_p, the average value of the CI's relative precision (CI HL over absolute value of the point estimate) as a percentage, and the estimated coverage of the CI (as a percentage), respectively. Columns 10 and 11 contain the average final batch size ($\overline{m}$) and average final batch count ($\overline{b}$), respectively. Finally, Columns 12 and 13 display the standard deviation of the CI HL and the average number of truncated observations (from each replication), respectively.

An examination of Table 1 reveals that FIRQUEST performed well for $p \le 0.95$ with respect to average CI relative precision and coverage probability. For example, for $p = 0.9$ and replication length $n = 40{,}000$ ($R = 5$), FIRQUEST reported a CI coverage probability of 94.9%. However, for the extreme value of $p = 0.99$, FIRQUEST exhibited slight undercoverage when the total sample size was less than 500,000, as it reported estimated CI coverage probabilities lower than the nominal value of 95%. For example, for $p = 0.99$ FIRQUEST reported an estimated CI coverage probability of 90.8%. Notably, such small sample sizes are inadequate for computing reliable CIs for this extreme quantile (Lolos et al. 2024, Section 4.3). Overall, we believe that FIRQUEST handled this challenging process effectively.

5 CONCLUSIONS

In this paper, we presented FIRQUEST, a fully automated fixed-sample-size procedure for computing point estimators and CIs for steady-state quantiles based on independent replications. Initial experimentation with successive customer waiting times in an M/M/1 system showed that FIRQUEST met the CI coverage probability requirements even when it was provided with relatively small total sample sizes. Topics of future work are (i) performance evaluation of FIRQUEST with additional test processes, (ii) estimation of multiple quantiles, and (iii) expansion of the framework for estimating the conditional value at risk.

REFERENCES

Alexopoulos, C., J. H. Boone, D. Goldsman, A. Lolos, K. D. Dingeç, and J. R. Wilson. 2020. "Steady-State Quantile Estimation Using Standardized Time Series". In *Proceedings of the 2020 Winter Simulation Conference*, 289–300 https://doi.org/https://doi.org/10.1109/WSC48552.2020.9384130.

Alexopoulos, C., and D. Goldsman. 2004. "To Batch or Not to Batch?". *ACM Transactions on Modeling and Computer Simulation* 14(1):76–114.

Alexopoulos, C., D. Goldsman, A. Lolos, K. D. Dingeç, and J. R. Wilson. 2025. "Steady-State Quantile Estimation Using Standardized Time Series". Technical report, Gebze Technical University, Georgia Institute of Tecnology, and North Carolina State University. https://www2.isye.gatech.edu/christos/papers/stsmsr5.pdf, Accessed 30th April.

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.

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*, 1814–1825 https://doi.org/https://doi.org/10.1109/WSC.2018.8632359.

Asmussen, S., and P. W. Glynn. 2007. *Stochastic Simulation: Algorithms and Analysis*. New York: Springer Science+Business Media.

Billingsley, P. 1999. *Convergence of Probability Measures*. 2nd ed. New York: John Wiley & Sons.

Dingeç, K. D., C. Alexopoulos, D. Goldsman, A. Lolos, and J. R. Wilson. 2022. "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, 111–130. New York: Springer.

Dingeç, K. D., C. Alexopoulos, D. Goldsman, A. Lolos, and J. R. Wilson. 2025. "Geometric-Moment Contraction, Stationary Processes, and Their Indicator Processes, I: Theory". Technical report, Gebze Technical University.

Fishman, G. S. 2001. *Discrete-Event Simulation: Modeling, Programming, and Analysis*. New York: Springer-Verlag.

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.

Goldsman, D., M. S. Meketon, and L. W. Schruben. 1990. "Properties of Standardized Time Series Weighted Area Variance Estimators". *Management Science* 36:602–612.

Law, A. M. 2015. *Simulation Modeling and Analysis*. 5th ed. New York: McGraw-Hill.

Lolos, A. 2023. *Effective Estimation of Marginal Quantiles in Steady-State Simulations*. Ph. D. thesis, H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0205, USA.

Lolos, A., C. Alexopoulos, D. Goldsman, K. D. Dingeç, A. Mokashi, and J. R. Wilson. 2023. "A Fixed-Sample-Size Method for Estimating Steady-State Quantiles". In *Proceedings of the 2023 Winter Simulation Conference*, 457–468 https://doi.org/https://doi.org/10.1109/WSC60868.2023.10407872.

Lolos, A., C. Alexopoulos, D. Goldsman, K. D. Dingeç, A. Mokashi, and J. R. Wilson. 2025. "FQUEST: A Fixed-Sample-Size Procedure for Estimating Steady-State Quantiles". Technical report, Georgia Institute of Technology.

Lolos, A., J. H. Boone, C. Alexopoulos, D. Goldsman, K. D. Dingeç, A. Mokashi *et al*. 2022. "A Sequential Method for Estimating Steady-State Quantiles Using Standardized Time Series". In *Proceedings of the 2022 Winter Simulation Conference*, 73–84 https://doi.org/https://doi.org/10.1109/WSC57314.2022.10015283.

Lolos, A., J. H. Boone, C. Alexopoulos, D. Goldsman, K. D. Dingeç, A. Mokashi *et al*. 2024. "SQSTS: A Sequential Procedure for Estimating Steady-State Quantiles Using Standardized Time Series". *Journal of Simulation* 18:988–1010 https://doi.org/https://doi.org/10.1080/17477778.2024.2362438.

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". *The 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 Principal Data Scientist in the Business Intelligence and Data Science team within the Membership and Deposit Products Division at Navy Federal Credit Union. He received his PhD in Operations Research from the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology. His email address is thnlolos@gmail.com.

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 served as *Proceedings* Co-Editor of the Winter Simulation Conference in 1995, Associate Program Chair in 2006, and on the Board of Directors of WSC between 2008–2016. 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/users/christos-alexopoulos.

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 DİNGEÇ 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 kdingec@yahoo.com.

ANUP C. MOKASHI is a Principal Data Scientist at Memorial Sloan Kettering Cancer Center. 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.