Proceedings of the 2019 Winter Simulation Conference
N. Mustafee, K.-H.G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, eds.

USING SIMULATION TO APPROXIMATE THE MINIMUM COST OF A FINITE SET OF ALTERNATIVES

Cuicui Zheng

Department of Computer Science
New Jersey Institute of Technology
323 Dr Martin Luther King Jr Blvd
Newark, NJ 07102, USA

James Calvin

Department of Computer Science
New Jersey Institute of Technology
323 Dr Martin Luther King Jr Blvd
Newark, NJ 07102, USA

ABSTRACT

We consider the problem of approximating the minimum cost of a finite set of alternative systems. We can not directly observe the cost of the systems, but we can estimate the cost using simulation. The simulation run lengths are adaptively chosen for each system. We describe an optimization algorithm and establish a bound on the error convergence rate. Compared with a single system, the error grows by an additional factor of the square root of the logarithm of the number of systems and the simulation budget.

1 INTRODUCTION

Consider the following optimization problem. We are interested in choosing a parameter from a finite set that optimizes the long-run average performance of a stochastic system. The system is too complicated to allow for analytic treatment, but we can simulate the performance at each parameter value.

Suppose that the parameter set is $\Theta = \{\theta_1, \theta_2, \ldots, \theta_\nu\}$, and the corresponding performance values are $\{\mu_1, \mu_2, \ldots, \mu_\nu\}$. For simplicity assume that the $\{\mu_i\}$ are distinct, and without loss of generality assume that $\mu_1 < \mu_2 < \cdots < \mu_\nu$. For each θ_i we can run a simulation, observing i.i.d. random variables $\{Y_{i,j} : j = 1, 2, \cdots\}$ with mean μ_i and variance σ_i^2. We assume that the $\{Y_{i,j}, 1 \le i \le \nu, j \ge 1\}$ are mutually independent and defined on a common probability space $(\Omega, \mathscr{F}, P)$. We further assume that for some $\varepsilon > 0$,

$$EY_{i,j}^{2+\varepsilon} < \infty.$$

We can adaptively increase the simulation lengths n_i so that the simulation effort is concentrated on the parameter values that appear most promising over time. Let n_i denote the number of simulation iterations at the ith system; that is, we have computed estimates based on $\{Y_{i,1}, \ldots, Y_{i,n_i}\}$. Let

$$n = \sum_{i=1}^{\nu} n_i.$$

Let $\mu_{n,i}$ denote the sample mean and $\sigma_{n,i}^2$ the sample variance for the simulation of the ith system:

$$\mu_{n,i} = \frac{1}{n_i} \sum_{j=1}^{n_i} Y_{i,j}$$

and

$$\sigma_{n,i}^2 = \frac{1}{n_i - 1} \sum_{j=1}^{n_i} \left(Y_{i,j} - \mu_{n,i}\right)^2.$$

After a total of n simulation steps, we have estimates $\mu_{n,i}$ and $\sigma_{n,i}$ such that

$$Z_{n,i} \equiv \frac{\sqrt{n_i}}{\sigma_{n,i}} (\mu_{n,i} - \mu_i) \xrightarrow{d} N(0,1) \tag{1}$$

as $n_i \to \infty$, for $i = 1, 2 \ldots, \nu$. We use the notation $X_n \xrightarrow{d} F$ to indicate that the sequence of random variables $\{X_n\}$ converges in distribution to the distribution F, and $X_n \xrightarrow{P} X$ to denote that the random variables $\{X_n\}$ converge in probability to the random variable X. We denote the standard normal distribution by $N(0,1)$, and its cumulative distribution function by Φ.

Let $\mu_{n,*} = \min_{1 \le i \le \nu} \mu_{n,i}$ and let i_n^* be the corresponding index, so that $\mu_{n,*} = \mu_{n,i_n^*}$; this will serve as our estimate of $\mu_1 = \min_i \mu_i$. We will construct an algorithm for running the simulations to efficiently estimate μ_1. Our goal is to construct a small interval $[\alpha_n, \beta_n]$ such that

$$P(\alpha_n \le \mu_1 \le \beta_n) \to 1$$

as $n \to \infty$. If we knew in advance that the first system was the best, then we could allocate all observations to that system and the central limit theorem implies that

$$P\left(\mu_{n,1} - \frac{\sigma_{n,i}}{\sqrt{n}} z_\alpha \le \mu_1 \le \mu_{n,1} + \frac{\sigma_{n,i}}{\sqrt{n}} z_\alpha\right) \to 1 - \alpha$$

as $n \to \infty$, where

$$1 - \Phi(z_\alpha) = \frac{\alpha}{2}.$$

Therefore, if γ_n is any increasing sequence going to $+\infty$, then the probability that μ_1 is contained in an interval of half-width

$$\frac{\sigma_{n,i}}{\sqrt{n}} \gamma_n \tag{2}$$

tends to 1. Without such advance knowledge, we must allocate observations to systems $2, 3, \ldots, \nu$, and so our enclosing interval should be larger; that is, it implies a lower bound on the growth of γ_n. Such a lower bound is obtained in Theorem 1 below.

2 BACKGROUND

To decide the best system using minimum cost in finite alternative systems by stochastic simulation is the problem we focused on. (Kim and Nelson 2001) developed procedures for selecting the best or near-best of a finite number of simulated systems and compared with indifference-zone procedures. (Ma and Henderson 2019) discussed similar problems as predicting simulation budget and guarantee probably approximately correct selection.

Approximating the minimum cost value is related to the problem of selecting the system with the smallest cost value. The latter problem is reviewed in (Goldsman et al. 2002; Lee and Nelson 2016). Our main reason for concentrating on the problem of approximating the minimum value instead of the minimizing parameter is that selecting the best system makes more sense for small, or at least finite, sets of alternatives. In cases of a continuum of alternatives, possibly with many isolated global minimizers, selecting the best system is not as well defined as approximating the minimum value. Nevertheless, the two problems are clearly related.

We take as given the performance measures $\{\mu_i\}$. An alternative is to start with a prior probability distribution on the $\{\mu_i\}$. This approach of Bayesian optimization is surveyed in Chick (2000) and Frazier (2018).

Our approach might be called a single-stage approach, in that the algorithm does not perform preliminary simulations that are used to plan subsequent simulations. Such single-stage procedures are presented in

Nakayama (2009), where the emphasis is on constructing asymptotically valid confidence intervals for the difference of cost values. We do not construct confidence intervals in the sense of that paper.

The algorithm described in this paper is similar to optimization algorithms described in works on continuous Bayesian optimization; for example see Calvin et al. (2017). In that work, the goal was to minimize a continuous function defined on the unit interval. The unknown function was assumed to be a sample path of a Wiener process. The algorithm adaptively chooses points to evaluate the random function f, with the error $\Delta_n(f)$ after n evaluations taken to be the difference between the smallest observed function value and the global minimum. It was shown that for all $r \in [1,\infty)$ and for all $p \in [1,\infty)$ there is a version of the algorithm (depending on those quantities) such that

$$(E|\Delta_n(f)|^p)^{1/p} \leq c \cdot n^{-r}$$

for a constant c and for all n. That is, any polynomial error rate can be obtained. This compares with the optimal nonadaptive error rate of $n^{-1/2}$.

The main motivation for our paper is to obtain insight into the discrete optimization problem as the number of alternatives $\nu \to \infty$. In order to formulate a reasonable version of this question we impose some structure on the $\{\mu_i\}$. In Section 5 we consider a smooth cost function $f : [0,1] \to \mathbb{R}$, that can only be evaluated by simulation. We adaptively run simulations at $f(i/\nu)$, $1 \leq i \leq \nu$, according the the proposed algorithm. As $\nu \to \infty$, we show that the rate at which the evaluations concentrate on the minimizer depends on f through the second derivative at the minimizer.

3 THE ALGORITHM

Recall that i_n^* is the index of the system with the smallest estimate $\mu_{n,i}$ after n simulation steps.

Define

$$g_{n,\nu} \equiv \sqrt{2\log(n\nu)}\frac{\sigma_{n,i_n^*}}{\sqrt{n_{i_n^*}}}$$

and

$$\rho_i^n \equiv \frac{\sigma_{n,i}^2}{n_i}\frac{1}{\left(\mu_{n,i} - \mu_{n,i_n^*} + g_{n,\nu}\right)^2}.$$

The idea of the algorithm is, at step n, to simulate the system i with the largest ρ_i^n. Roughly, this can be thought of as maximizing the probability that the next evaluation is below $\mu_{n,i_n^*} - g_{n,\nu}$. The "gap" $g_{n,\nu}$ is chosen large enough that these probabilities (and therefore the ρ_i^n) go to zero.

To simplify the exposition, we will assume that each system is initially simulated for two steps so that the sample mean and variances are defined. The algorithm for a simulation budget $N > 2\nu$ follows.

1. Simulate 2 steps for each of the ν systems, set $n_i = 2$, $1 \leq i \leq \nu$, and compute $\mu_{n,i}$ and $\sigma_{n,i}$, $1 \leq i \leq \nu$. Set i_n^* to the (first) index with minimal sample mean, and set $n = 2\nu$.
2. Compute ρ_i^n for each $i \leq \nu$, and let k be the (first) index with $\rho_k^n \geq \rho_i^n \ \forall i$.
3. Simulate the kth system for one step, and update the sample mean and variance of the kth system.
4. Set $n_k = n_k + 1$, $n = n + 1$, and if $n < N$ return to step 2.

We will only consider times at which we are about to evaluate the currently best system; that is, when $\rho_{i_n^*}^n \geq \rho_k^n$ for all k. Note that at such a time

$$\rho_k^n \leq \rho_{i_n^*}^n \leq \frac{1}{2\log(n\nu)}; \tag{3}$$

that is, for each i,

$$\frac{\sigma_{n,i}^2}{n_i}\frac{1}{\left(\mu_{n,i} - \mu_{n,i_n^*} + g_{n,\nu}\right)^2} \leq \frac{1}{2\log(n\nu)}.$$

The event that the cost of the ith system is above our lower bound $\mu_{n,i_n^*} - g_{n,\nu}$ is

$$\begin{aligned}
\{\mu_i > \mu_{n,i_n^*} - g_{n,\nu}\} &= \left\{\mu_{n,i} - \frac{\sigma_{n,i}}{\sqrt{n_i}} Z_{n,i} > \mu_{n,i_n^*} - g_{n,\nu}\right\} \quad \text{by (1)} \\
&= \left\{Z_{n,i} < \frac{\sqrt{n_i}}{\sigma_{n,i}} \left(\mu_{n,i} - \mu_{n,i_n^*} + g_{n,\nu}\right)\right\} \\
&= \left\{Z_{n,i} < \frac{1}{\sqrt{\rho_i^n}}\right\} \\
&\supset \left\{Z_{n,i} < \sqrt{2\log(n\nu)}\right\} \quad \text{by (3)} \\
&\supset \left\{Z_{n,i} < \sqrt{2\log(n_i\nu)}\right\}.
\end{aligned}$$

Therefore, by (1),

$$P\left(\mu_i > \mu_{n,i_n^*} - g_{n,\nu}\right) \to 1$$

as $n \to \infty$. It follows that

$$\begin{aligned}
P(\Delta_n \le g_{n,\nu}) &= P(\mu_{n,i_n^*} - \mu_1 \le g_{n,\nu}) \\
&= P\left(\bigcap_{i=1}^{\nu} \{\mu_i > \mu_{n,i_n^*} - g_{n,\nu}\}\right) \\
&\to 1
\end{aligned}$$

as $n \to \infty$.

The amount of simulation of the ith system satisfies

$$n_i \approx \frac{2\log(n\nu)\sigma_{n,i}^2}{(\mu_{n,i} - \mu_{n,i_n^*} + g_{n,\nu})^2},$$

and so for $i > 1$,

$$\frac{n_i}{\log(n\nu)} \xrightarrow{P} \frac{2\sigma_i^2}{(\mu_i - \mu_1)^2}$$

as $n \to \infty$. Also,

$$\frac{n_{i_n^*}}{n} \xrightarrow{P} 1.$$

This implies (1) ((Billingsley 1968), Theorem 17.1).

4 MAIN RESULT

Set

$$\mathscr{H} = \sum_{i=2}^{\nu} \frac{\sigma_i^2}{(\mu_i - \mu_1)^2}.$$

This is a measure of how hard the optimization is. If the small μ_i's are clustered around the minimum, then the algorithm needs to spread out the effort over multiple systems that seem promising; this corresponds to a large $\mathscr{H}$. If μ_2 is much larger than the minimum μ_1, then the search can concentrate on μ_1 and $\mathscr{H}$ is small. Then we have

$$n_{i_n^*} \approx n - 2\log(n\nu)\mathscr{H}$$

and

$$\frac{n_{i_n^*}}{n} \xrightarrow{P} 1.$$

The following limit theorem shows how small our interval containing the minimizer with probability approaching one can be.

Theorem 1 Consider a system with ν alternatives. Let

$$\Delta_n \equiv \mu_{n,i_n^*} - \mu_1$$

denote the approximation error after n simulation steps. Then

$$\lim_{n\to\infty} P\left(-\frac{\sigma_1}{\sqrt{n}} \frac{\sqrt{2\log(n\nu)}}{\sqrt{1-2\mathscr{H}\log(n\nu)/n}} \le \Delta_n \le \frac{\sigma_1}{\sqrt{n}} \frac{\sqrt{2\log(n\nu)}}{\sqrt{1-2\mathscr{H}\log(n\nu)/n}} \right) = 1.$$

This shows the slowdown we suffer compared to (2). We see that whereas the sequence γ_n introduced at (2) can grow arbitrarily slowly, for our algorithm it must grow at rate $\sqrt{\log(n\nu)}$.

5 CONCENTRATION RATE

The $\{\mu_i\}$, and hence $\mathscr{H}$, could be quite arbitrary. In this section we consider the case where the goal is to approximate the minimum of a continuous cost function by simulating values at a fixed grid of points. Denote the cost function by $f : [0,1] \to \mathbb{R}$. We estimate the values $f(i/\nu)$, $i = 1,\ldots,\nu$. With our previous notation, $\{\mu_i, 1 \le i \le \nu\} = \{f(i/\nu), 1 \le i \le \nu\}$, though the ordering is different in general.

In this section, assume that $\sigma_i \equiv \sigma$ for each $1 \le i \le \nu$. Let us suppose that $f \in C^2([0,1])$ with unique global minimizer $t^* \in (0,1)$ with $f''(t^*) > 0$. Let Π_n denote the proportion of evaluations that are not at the current estimator of the minimizer:

$$\Pi_n = 1 - \frac{n_{i_n^*}}{n}.$$

Set

$$B(f,\nu,n) \equiv \frac{\pi}{4}\nu \frac{\sqrt{\sigma}}{\sqrt{\beta(f,\nu)}} \left(\frac{2\log(n\nu)}{n}\right)^{1/4},$$

where $\beta(f,\nu) \to f''(t^*)$ as $\nu \to \infty$.

Theorem 2 As the number of evaluations n tends to infinity,

$$P(\Pi_n \le B(f,\nu,n)) \to 1.$$

The bound $B(f,\nu,n)$ increases roughly linearly in the discretization order ν, and as the square root of the noise standard deviation σ. For large ν, $\beta(f,\nu) \approx f''(t^*)$ and so $B(f,\nu,n)$ increases as $1/\sqrt{f''(t^*)}$. For f that increases rapidly moving away from the global minimizer, i.e., for large $f''(t^*)$, the algorithm can concentrate the evaluations more near the minimizer, thus reducing Π_n.

6 NUMERICAL EXPERIMENTS

We did some experiments for the algorithm applied to the discretization of a smooth cost function, as described in the last section. The cost function is

$$f(x) = \frac{1}{2} + 2x - \cos(12x-9), \quad 0 \le x \le 1, \tag{4}$$

which is shown in Figure 1a. Figure 1b shows the discrete version of the optimization problem, where the ith system cost is $f(i/\nu)$ for $\nu = 30$ and $1 \le i \le \nu$.

Figure 1: Continuous and discretized cost functions.

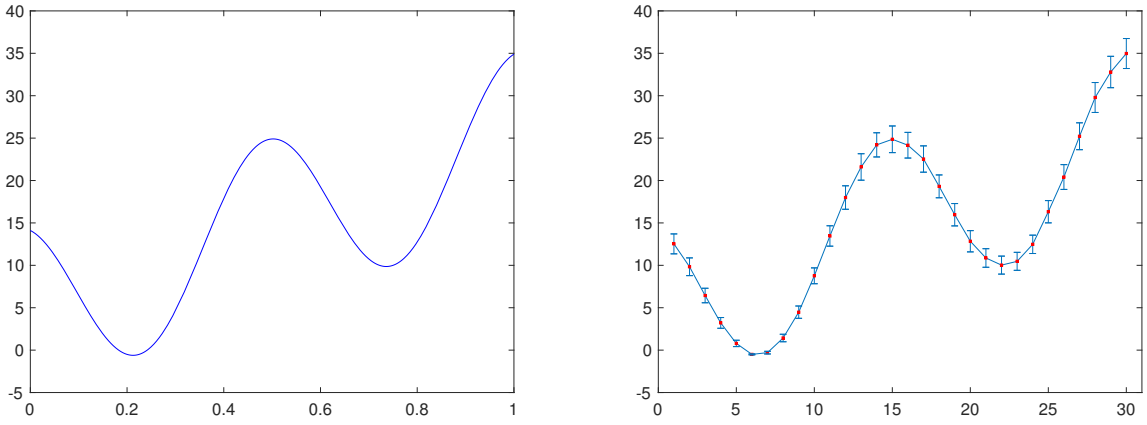

(a) Continuous cost function (4).

(b) Estimates of discretized cost function (4).

For a fixed number ν, our goal is to approximate the minimum of $f(i/\nu)$, $1 \leq i \leq \nu$. This will be based on evaluations

$$Y_{i,j} = f(i/\nu) + \xi_{i,j},$$

where the $\xi_{i,j}$ are independent standard normal random variables. Figure 1b shows the results of running the algorithm on the function shown in Figure 1a with $\nu = 30$ and $n = 40{,}000$. The error bars around each estimated function value indicate the standard deviation of the estimator. In Figure 2 we plot the normalized error for the algorithm applied to discretizations of $\nu = 50$ in the left plot and $\nu = 100$ in the right plot. From Theorem 1, the normalized error

$$\frac{\sqrt{n}}{\sigma_1 \sqrt{2\log(n\nu)}}$$

will lie in the interval $[-1,1]$ with probability approaching 1 as $n \to \infty$. The normalized error is plotted as a function of the iteration number n for $\nu = 50$ (left plot) and $\nu = 100$ (right plot). In Figure 3, we plot the proportion of evaluations made at points other than the current estimated minimizer, Π_n, divided by our upper bound $B(f,\nu,n)$ from Theorem 2. We evaluate the rate $\Pi_n/B(f,\nu,n)$ with $\nu = 50$ and $\nu = 100$ with $\sigma = 1$. The total iteration number is $40{,}000$. Theorem 2 implies that the ration should be below 1 with probability approaching 1 as $n \to \infty$. The results are plotted for $\nu = 50$ (left plot) and $\nu = 100$ (right plot).

7 CONCLUSIONS

We have constructed a single-stage procedure for adaptively controlling the simulation of multiple systems with the aim of efficiently approximating the minimum cost measure. We constructed an interval that contains the minimum cost measure with probability approaching 1 as the simulation budget grows. Compared with the ideal situation of a single system, the size of the enclosing interval grows by an additional factor of the square root of the logarithm of the number of systems and the simulation budget.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under Grant No. CMMI-1562466.

Figure 2: Normalized proportion of evaluations at sub optimal points.

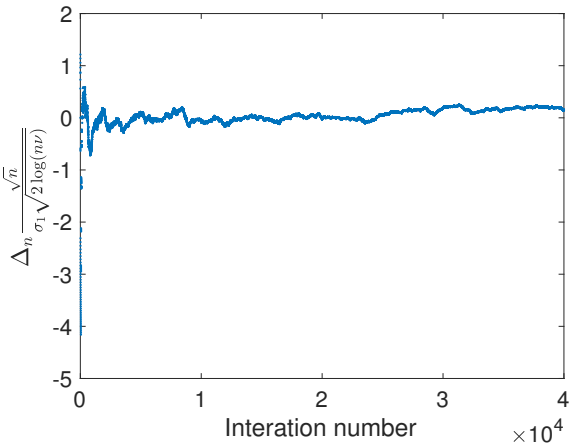

(a) Normalized proportion of evaluations at sub optimal points, $\nu = 50$.

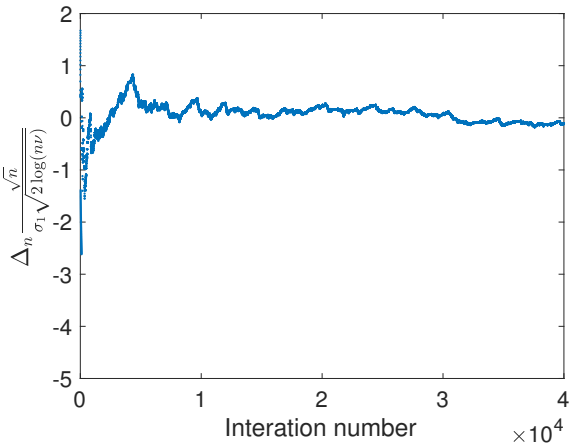

(b) Normalized proportion of evaluations at sub optimal points, $\nu = 100$.

Figure 3: Normalized proportion for increasing number of evaluations.

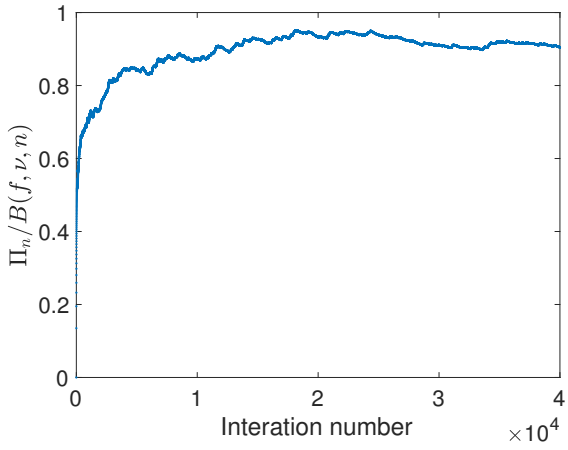

(a) Normalized error, $\nu = 50$.

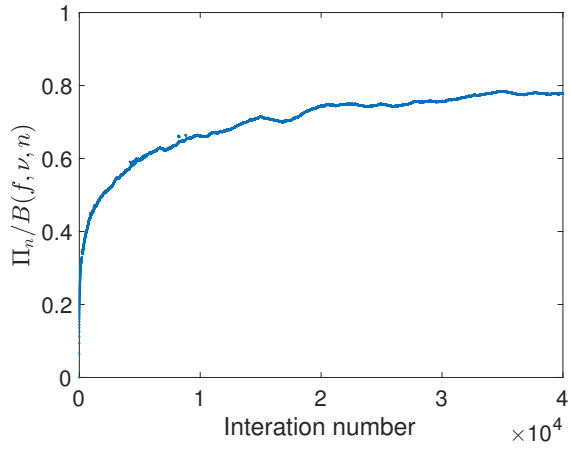

(b) Normalized error, $\nu = 100$.

REFERENCES

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

Calvin, J. M., M. Hefter, and A. Herzwurm. 2017. "Adaptive Approximation of the Minimum of Brownian Motion". *Journal of Complexity* 39:17–37.

Chick, S. E. 2000. "Bayesian Methods: Bayesian Methods for Simulation". In *Proceedings of the 2000 Winter Simulation Conference*, edited by K. K. J. A. Joines, R. R. Barton and e. P. A. Fishwick, 109–118. Orlando, Florida: Society for Computer Simulation International.

Frazier, P. I. 2018. "A Tutorial on Bayesian Optimization". *eprint arXiv* 1807.02811.

Goldsman, D., S. Kim, W. S. Marshall, and B. L. Nelson. 2002. "Ranking and Selection for Steady-State Simulation: Procedures and Perspectives". *Informs Journal on Computing* 14(1):2–19.

Kim, S.-H., and B. L. Nelson. 2001. "A Fully Sequential Procedure for Indifference-zone Selection in Simulation". *ACM Transactions on Modeling and Computer Simulation* 11(3):251–273.

Lee, S., and B. L. Nelson. 2016. "General-purpose Ranking and Selection for Computer Simulation". *IIE Transactions* 48(6):555–564.

Ma, S., and S. G. Henderson. 2019. "Predicting the Simulation Budget in Ranking and Selection Procedures". *ACM Transactions on Modeling and Computer Simulation* 29(3):14:1–14:25.

Nakayama, M. K. 2009. "Asymptotically Valid Single-stage Multiple-comparison Procedures". *Journal of Statistical Planning and Inference* 139(4):1348–1356.

AUTHOR BIOGRAPHIES

CUICUI ZHENG is a PhD candidate in the Department of Computer Science in the Ying Wu College of Computing at the New Jersey Institute of Technology. She finished her bachelor's degree in Information Security at Northeastern University, Shenyang China, 2015. Her research interests include global optimization and its application to image registration and clustering. Her email address is cz296@njit.edu.

JAMES CALVIN is a Professor in the Department of Computer Science in the Ying Wu College of Computing at the New Jersey Institute of Technology. Besides simulation modeling and analysis, his research interests focus on developing global optimization algorithms and establishing complexity bounds for global optimization problems. His email address is calvin@njit.edu.