

USING QUANTILES IN RANKING AND SELECTION PROCEDURES

Jennifer M. Bekki
John W. Fowler
Gerald T. Mackulak

Industrial Engineering Department
Arizona State University
Tempe, AZ 85287-5906, U.S.A.

Barry L. Nelson

Dept. of Industrial Eng. & Management Sciences
Northwestern University
Evanston, IL 60208-3119, U.S.A.

ABSTRACT

A useful performance measure on which to compare manufacturing systems is a quantile of the cycle time distribution. Unfortunately, aside from order statistic estimates, which can require significant data storage, the distribution of quantile estimates has not been shown to be normally distributed, violating a common assumption amongst ranking-and-selection (R&S) procedures. To address this, we provide empirical evidence supporting an approach using the mean of a group of quantile estimates as the comparison measure. The approach is detailed and illustrated through experimentation on four M/M/1 queues in which the 0.9 cycle-time quantile is the performance measure. Results in terms of simulation effort and accuracy are reported and compared to results obtained using the macro-replications approach for inducing normality as well as to results obtained by applying R&S procedures to quantile estimates directly. The suggested procedure is shown to provide significant savings in simulation effort while sacrificing very little in accuracy.

1 INTRODUCTION

Ranking and Selection (R&S) procedures are statistical techniques designed to select the best system, or a subset of systems containing the best, from a group of alternatives. The procedures guarantee a user specified probability of selecting the best system, and typical assumptions include that the response measure is normally distributed and that the analysis is based on expected values of a single stochastic response. The variance of the response at different treatments may be equal or unequal, and the variance at treatments may be known or unknown. Independence between responses is generally assumed, but techniques exist which employ common random numbers (CRNs), intentionally inducing dependence between treatments, but often reducing required sample sizes. An example of such a technique can be found in Nelson and Matejcek (1995).

The simplest and most common application of R&S approaches involves identifying the best system from a relatively small (< 20) number of alternatives based on mean performance measures such as time in system (for larger numbers of systems, approaches such as that suggested in Nelson et al. (2001) should be applied). However, another useful measure on which to make comparisons is a specific quantile of the cycle-time distribution. Comparisons based on a quantile of the cycle-time distribution allow decision makers to determine the best system (where a system with a smaller quantile is considered better) from a group of competing scenarios at various levels of risk. Conservative decision makers could use the 0.9 cycle-time quantile on which to base their system selection, indicating that they have chosen the best system as the system in which they will meet 90% of their demand on time in the shortest amount of time. Or, conversely, more aggressive decision makers could base the selection on the 0.7 cycle-time quantile, being comfortable with the possibility of quoting lead times incorrectly 30% of the time. In many cases, these comparisons produce results different than those obtained by simply comparing the means of the cycle-time distribution, even when the expected value of the cycle-time distribution between the systems was the same. For instance, two normal distributions with the same mean but different standard deviations produce two different 0.8 quantile estimates. As a result, comparisons based on cycle-time quantiles provide an enhanced level of detail for decision making over simple comparisons on the mean.

Unfortunately, while quantile estimates based on order statistics satisfy the Central Limit Theorem (David, 1981), they traditionally require the storage of all observations, making their use less attractive. Indirect quantile estimates, on the other hand, which are more attractive in terms of data storage, have not been shown to be normally distributed, violating the assumption of the R&S procedure that the response measure be normally distributed. One solution to this problem is to develop a new R&S procedure that directly

compares raw quantile estimates of competing systems and which does not require that the estimates be normally distributed. Alternatively, a method could be developed to utilize already developed R&S procedures by transforming the quantile estimates to be more normally distributed.

One such solution is suggested by Goldsman et al. (1991). They suggest making macro-replications, each consisting of multiple micro-replications. The average of the values of the performance measure from the micro-replications are then taken as the single estimate of the performance measure from the macro-replication. The performance measure from the macro-replication, then, represents a single data point to be used as part of a R&S procedure. Based on the Central Limit Theorem (CLT), the averages of these data points from the micro-replications will be more nearly normally distributed, and the underlying assumptions of the procedures will not be violated.

The drawback of this approach is that the use of macro-replications requires a large amount of simulation data. For example, if an initial sample size of five data points was desired for a R&S procedure, 5 macro-replications would be required for each system. If the number of micro-replications making up each macro-replications was 20, then the total number of initial observations of the performance measure from each system would be 100 ($= 20 \times 5$). If the R&S procedure then dictated additional sampling, each additional sample for each system would again represent the average of 20 micro-replications. The likelihood of additional sampling increases as the variability of the performance measure increases. Therefore, as the variability in the performance measure increases, the simulation effort can quickly explode, making the implementation of R&S in this manner impractical, particularly for large models such as those found in the semiconductor industry in which run lengths can be extremely long.

Ideally, an approach for using quantile estimates as the performance measure would allow existing R&S procedures to be used while not requiring an unreasonable sample size. The contribution of such an approach would not be in altering the R&S procedures themselves. These procedures are generic to the performance measure on which comparisons are made, provided that the distribution of the performance measure estimator meets certain assumptions, and simply using a cycle-time quantile as the performance measure over the mean cycle-time, for instance, would not provide anything novel. However, given that quantiles estimates are not normally distributed and that most R&S procedures require a normally distributed performance measure, a procedure for transforming quantile estimates into a representative normal distribution while not requiring enormous simulation effort would be valuable.

The purpose of this paper is to provide empirical evidence supporting an approach for using quantile estimates as the comparison measure for a single-stage R&S procedure

for comparing systems with unequal but known variances. The procedure uses the variance of an initial set of quantile estimates to determine the corresponding variance of the distribution surrounding the mean of a group of quantile estimates. The distribution surrounding the group mean is, by the CLT, more normally distributed than the original quantile estimates, and the variance of these means are used with the R&S procedure to determine the required number of observations (group means) from each system. The approach capitalizes on the initial set of runs used to determine the variance of the quantile estimates by reusing the same runs toward observations for the R&S procedure. Once the initial runs are done, however, the procedure is very similar to the macro-replications approach, as it requires a group of runs to yield a single observation for use in a R&S procedure. Consequently, the approach is particularly useful for comparisons between systems in which the variance of the comparison metric is not high but where the effort to obtain a single estimate of that performance measure is very high. In such cases, the majority of simulation effort comes from initial data collection, which is efficient in the proposed procedure. An example of such a system is a model of a semiconductor wafer fab. Typically, these models take a very long time to run, but usually do not produce highly variable results.

The remainder of the paper is organized as follows. Section 2 details the proposed procedure for the use of quantiles as a comparison measure with R&S procedures, and Section 3 then gives experimentation using the procedure and discusses the results. Conclusions and areas for future work are then given in Section 4.

2 PROCEDURE FOR USING QUANTILES IN R&S

The objective of the procedure is to use quantile estimates in a R&S procedure, which assumes that the variable upon which comparisons are made is normally distributed. To do so, we use the distribution surrounding the mean of a group of quantile estimators, which we know to be closer to normal than the original quantile estimates. Moreover, from basic probability, a relationship exists between the original set of quantile estimates, $\hat{y}_q(1), \hat{y}_q(2), \dots, \hat{y}_q(n)$ and the distribution surrounding their mean, $\hat{\bar{y}}_q$:

$$E[\hat{\bar{y}}_q] = E[\hat{y}_q] \text{ and } Var[\hat{\bar{y}}_q] = Var[\hat{y}_q]/n.$$

Therefore, assuming we know (or can estimate very well), the variance of the original quantile estimates, we can also determine the variance of the distribution surrounding the mean of a group of n quantile estimates. The R&S procedure is then performed on the $\hat{\bar{y}}_q$ values, which represent samples from a new random variable that shares properties with the original random variable of interest, but which

also has the desired property of coming from a normal distribution. A detailed description of the procedure follows:

Grouped Quantile Approach: Using Precise Estimates of Variance with Known Variance R&S Procedure

1. Make n_f independent simulation replications of each system $i = 1, 2, \dots, k$. We recommend $n_f \geq 50$.
2. From each simulation replication, obtain an estimate of the q quantile, $\hat{y}_q(1), \hat{y}_q(2), \dots, \hat{y}_q(n_f)$.
3. Using $\hat{y}_q(1), \hat{y}_q(2), \dots, \hat{y}_q(n_f)$, estimate $Var[\hat{q}]_i$ for each system $i = 1, 2, \dots, k$.
4. For each system, estimate the variance of the distribution surrounding the mean, \bar{y}_q , of a group of n_b quantile estimates, where $\hat{V}ar[\hat{y}_q]_i = \frac{Var[\hat{y}_q]_i}{n_b}$. Use $Var[\hat{q}]_i$ calculated in Step 3. We recommend $n_b \geq 20$.
5. Select confidence level P^* (PCS) and indifference zone parameter, $\delta^* > 0$.
6. Use the following formulas to determine the number of observations, n_i required from each system (Bechhofer et al., 1995) Use $\hat{V}ar[\hat{y}_q]_i$ in place of σ_i . Values for $Z_{1-1/2}^{(1-P^*)}$ can be found in Bechhofer et al. (1995).

$$n_i = \begin{cases} \left\lceil \left[\sigma_i \sqrt{2(\sigma_1^2 + \sigma_2^2)} (Z_{1,1/2}^{1-P^*} / \delta^*)^2 \right] \right\rceil & k = 2 \\ \left\lceil 2 \left(\sigma_i Z_{1-1/2}^{(1-P^*)} / \delta^* \right)^2 \right\rceil & \text{o.w.} \end{cases}$$

7. Make $n_t = (n_i * n_b) - n_f$ independent simulation replications of each system $i = 1, 2, \dots, k$.
8. From each simulation replication, obtain an estimate of the q quantile, $\hat{y}_q(1), \hat{y}_q(2), \dots, \hat{y}_q(n_t)$.
9. Select the system with the smallest $\bar{y}_{q_i}(n_t)$, where $\bar{y}_{q_i}(n_t)$ represents the mean of all observations obtained from each system, including the those obtained from the original n_f simulation runs.

3 EMPIRICAL EVALUATION

To evaluate the approach described in Section 2, it was applied to four M/M/1 systems. The performance measure upon which comparisons between systems were made was the 0.9 CT quantile, where the best system was chosen as the one with the smallest 0.9 cycle-time quantile. The four M/M/1 systems were configured so that systems 1, 3, 4 have true 0.9 cycle-time quantile values which are exactly δ greater than the true value of the 0.9 CT quantile for system 2. This configuration is known to be the most difficult in

which to identify the best system. Table 1 gives the arrival and service rates selected for the four M/M/1 systems.

Using the four systems described in Table 1, the approach based on using a precise estimate of the variance in place of the true variance of the means, described in Section 2, was applied to select the best system. For comparison purposes, the macro-replications approach by Goldsman et al. (1991) was also used to select the best of the same five systems. To combine the macro-replications procedure with a ranking-and-selection procedure, a R&S procedure that assumes that the variance of the performance measure is unknown must be used, as the variance of the observations from the macro-replications is not known in advance. We use Rinott’s procedure (Rinott 1978) for this purpose since it is one of the simplest and most well known of the multi-stage ranking-and-selection procedures. Rinott’s procedure makes the following assumptions: there are $k \geq 2$ systems, X_{ij} is the j^{th} independent observation from system i , $X_{ij} \sim N(\mu_i, \sigma_i^2)$, where μ_i and σ_i^2 are unknown, and the data across systems are independent. To implement the procedure, the following steps are taken, where PCS represents the probability of a correct selection, and δ represents the minimum detectable difference between competing systems.

Rinott’s Procedure

1. Select confidence level $1 - \alpha$ (PCS), indifference zone parameter, $\delta > 0$, and first stage sample size ($n_0 \geq 2$).
2. Obtain Rinott’s constant $h = h(n_0, k, 1 - \alpha)$ from Wilcox (1984).
3. Obtain n_0 observations $X_{ij}, j = 1, 2, \dots, n_0$ from each system $i = 1, 2, \dots, k$.
4. Compute the sample variance of data from system i :

$$s_i^2 = \frac{1}{n_0 - 1} \sum_{j=1}^{n_0} (X_{ij} - \bar{X}_i(n_0))^2$$

5. Compute the number of observations required from system i :

$$N_i = \max \left\{ n_0, \left\lceil \frac{h^2 s_i^2}{\delta^2} \right\rceil \right\}.$$

Table 1: Description of M/M/1 systems used to evaluate approaches for using quantiles with R&S procedures.

System	Service Rate	Arrival Rate	0.9 Quantile
1	1.000	1.101	25.00
2	2.000	2.463	24.50
3	1.500	1.740	25.00
4	2.000	2.452	25.00

6. If $n_0 \geq N_i$, stop and choose system with largest $\bar{X}_i(n_0)$. Otherwise, take $N_i - n_0$ additional observations from each system i for which $N_i > n_0$ and then select the system with the smallest $\bar{X}_i(N_i)$.

Then, combining the macro-replications approach with Rinott's procedure gives the following procedure:

Macro-Replications with Rinott's Procedure

1. Select confidence level $1 - \alpha$ (PCS), indifference-zone parameter, $\delta > 0$, and first-stage sample size ($n_o \geq 2$) for Rinott's procedure.
2. Make $n_f = n_o * n_g$ independent simulation replications of each system $i = 1, 2, \dots, k$. Each independent simulation run represents a micro-replication, and n_g represents the number of micro-replications in a single macro-replication.
3. From each simulation replication, obtain an estimate of the q quantile, $\hat{y}_q(1), \hat{y}_q(2), \dots, \hat{y}_q(n_f)$.
4. Calculate the average of every n_g estimates of $[\hat{y}_q]_i$ from each system i . Each average value, $[\bar{\hat{y}}_q]_i$ represents a macro-replication, and, by the CLT, the distribution surrounding $[\bar{\hat{y}}_q]_i$ is more nearly normally distributed.
5. Compute the sample variance of the $[\bar{\hat{y}}_q]_i$ values from each system i using the formula presented in Step 4 of Rinott's procedure.
6. Compute the number of $[\bar{\hat{y}}_q]_i$ observations required from system i , N_i , using the formula presented in Step 5 of Rinott's procedure.
7. If $n_0 \geq N_i$ for all systems, stop and choose system with smallest $[\bar{\hat{y}}_q]_i(n_0)$, where $[\bar{\hat{y}}_q]_i(n_0)$ represents the mean of the n_o $[\bar{\hat{y}}_q]_i$ observations for each system i .
8. Otherwise,
 - (a) Make $n_t = n_g * (N_i - n_0)$ additional independent simulation replications from each system i in which $N_i > n_0$.
 - (b) From each simulation replication, obtain an estimate of the q quantile, $\hat{y}_q(1), \hat{y}_q(2), \dots, \hat{y}_q(n_t)$.
 - (c) Select the system with the smallest $[\bar{\hat{y}}_q]_i(N_i)$, where where $[\bar{\hat{y}}_q]_i(N_i)$ represents the mean of all $[\bar{\hat{y}}_q]_i$ observations obtained from each system.

In addition to the macro-replications approach, the proposed procedure was also compared to approaches using the original quantile estimates, without any attempt to correct for their non-normality. The quantile estimates were used both in combination with the known variance R&S procedure, after estimating their variance from an initial set of n_f quantile estimates, and with Rinott's procedure, assuming

we do not know the variance in advance. The details of these approaches follow; note that the procedure using Rinott's procedure is identical to the macro-replications approach using $n_g = 1$.

Raw quantiles with Rinott's Procedure

1. Select confidence level $1 - \alpha$ (PCS), indifference-zone parameter, $\delta > 0$, and first-stage sample size ($n_o \geq 2$) for Rinott's procedure.
2. Make n_0 independent simulation replications of each system $i = 1, 2, \dots, k$.
3. From each simulation replication, obtain an estimate of the q quantile, $\hat{y}_q(1), \hat{y}_q(2), \dots, \hat{y}_q(n_f)$.
4. Compute the sample variance of the $[\hat{y}_q]_i$ values from each system i using the formula presented in Step 4 of Rinott's procedure.
5. Compute the number of $[\hat{y}_q]_i$ observations required from system i , N_i , using the formula presented in Step 5 of Rinott's procedure.
6. If $n_0 \geq N_i$ for all systems, stop and choose system with largest $[\hat{y}_q]_i(n_0)$, where $[\hat{y}_q]_i(n_0)$ represents the mean of the n_o $[\hat{y}_q]_i$ observations for each system i . Otherwise,
 - (a) Make $n_t = N_i - n_0$ additional independent simulation replications from each system i in which $N_i > n_0$.
 - (b) From each simulation replication, obtain an estimate of the q quantile, $\hat{y}_q(1), \hat{y}_q(2), \dots, \hat{y}_q(n_t)$.
 - (c) Select the system with the smallest $[\hat{y}_q]_i(N_i)$, where where $[\hat{y}_q]_i(N_i)$ represents the mean of all $[\hat{y}_q]_i$ observations obtained from each system.

Raw Quantiles Using Precise Estimates of Variance R&S procedure

1. Make n_f independent simulation replications of each system $i = 1, 2, \dots, k$.
2. From each simulation replication, obtain an estimate of the q quantile, $\hat{y}_q(1), \hat{y}_q(2), \dots, \hat{y}_q(n_f)$.
3. Using $\hat{y}_q(1), \hat{y}_q(2), \dots, \hat{y}_q(n_f)$, estimate $Var[\hat{q}]_i$ for each system $i = 1, 2, \dots, k$.
4. Select confidence level P^* (PCS) and indifference zone parameter, $\delta^* > 0$.
5. Use the following formulas to determine the number of observations, n_i required from each system (Bechhofer et al., 1995). Use $\hat{Var}[\hat{q}]_i$ in place of σ_i . Values for $Z_{1-P^*}^{(1-P^*)}$ can be found in Bechhofer

et al. (1995).

$$n_i = \begin{cases} \left\lceil \sigma_i \sqrt{2(\sigma_1^2 + \sigma_2^2)} (Z_{1,1/2}^{1-P^*} / \delta^*)^2 \right\rceil & k = 2 \\ \left\lceil 2 \left(\sigma_i Z_{i-1,1/2}^{(1-P^*)} / \delta^* \right)^2 \right\rceil & \text{o.w.} \end{cases}$$

6. Make $n_t = n_i - n_f$ independent simulation replications of each system $i = 1, 2, \dots, k$.
7. From each simulation replication, obtain an estimate of the q quantile, $\hat{y}_q(1), \hat{y}_q(2), \dots, \hat{y}_q(n_t)$.
8. Select the system with the smallest $\bar{\hat{y}}_{q_i}(n_i)$, where $\bar{\hat{y}}_{q_i}(n_i)$ represents the mean of all observations obtained from each system, including the those obtained from the original n_f simulation runs.

For all four approaches (macro-replications, raw quantiles using the known-variance approach, raw quantiles using Rinott's procedure, and the suggested approach based on assuming known variance based on an estimate from initial runs), to obtain the 0.9 cycle time quantile estimates, independent simulation runs of each system were made, and the approach suggested by Bekki et al. (2006) and McNeill et al. (2003), based on the Cornish-Fisher expansion, was used. The simulation run length, m , was set at 250k observations, with an additional run length of 10% truncated. In all cases, the Probability of Correct Selection (PCS) for the R&S procedure was selected to be 0.95, and the indifference-zone parameter was chosen to be $\delta = 0.5$. Each of the four approaches was applied to the four M/M/1 system 500 times. Finally, for the approaches in which the variance is determined based on an initial set of runs and is then assumed to be known, the n_f parameter was set to 50 and the n_b parameter used was 20. In cases where Rinott's procedure was used, the n_0 parameter was set to 5.

After applying each of the approaches, Figure 1 shows the average percentage of simulation effort allocated to each system across the 500 replications. The categories on the x-axis represent the different approaches combined with the R&S procedure to determine the best system. This chart illustrates that when grouping is involved, the allocation of simulation effort is less sensitive to the variability of quantile estimates. For example, a larger amount of effort is allocated to system 1, the system with the highest traffic intensity and, therefore, greatest variability in quantile estimates, in each of the techniques that do not involving grouping the quantile estimates. This effect may be due to the fact that the variance of the mean of a group is equivalent to the variance of the ungrouped values divided by the number of observations in the group (see Section 2, making the estimates based on the mean of grouped values less variable in general than the variance of raw observations. Overall, Figure 1 illustrates that both approaches using raw quantiles are very sensitive

in terms of simulation effort allocation to the variability differences between systems.

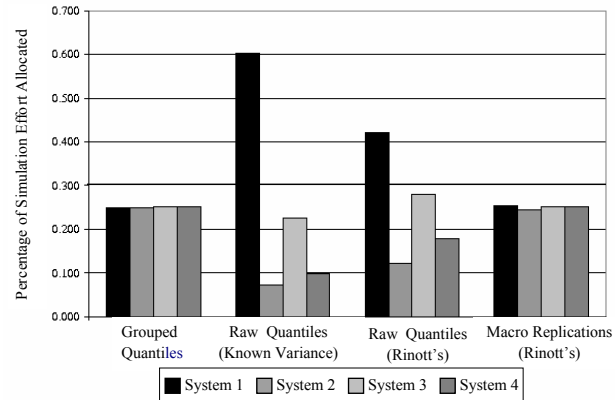


Figure 1: Allocation of simulation effort across the four M/M/1 systems.

Figure 2 gives a summary of the results in performance space. The y-axis of the chart shows the average simulation effort across all four systems, in terms of 250 thousands of observations, required for a single replication of R&S procedure using each of the four approaches. The x-axis shows a 0.95 confidence interval surrounding 1 - the realized PCS. The confidence intervals were built using the following formula, which utilizes the fact that the binomial distribution can be approximated with a normal distribution. In this formula, \hat{p} represents the proportion of times the correct system was chosen, $z_{1-\alpha/2}$ represents the $1 - \alpha/2$ quantile of a standard normal distribution, and n is the sample size of 500

$$\hat{p} \pm z_{1-\alpha/2} \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$

In Figure 2, approaches are comparatively better as they get closer to the lower left corner of the figure. In this corner, the realized PCS is high, and the simulation effort is low. Points that have the same simulation effort with lower realized PCS are weakly dominated, as are points that have the same realized PCS but require greater simulation effort. In Figure 2, the confidence interval surrounding the grouped quantiles approach overlaps with both the known and unknown variance R&S procedures done with raw quantiles. Consequently, in terms of the realized PCS, the grouped quantiles approach is statistically the same as both approaches based on raw quantiles. However, of the three approaches, the suggested grouped quantiles approach requires considerably less (25% - 35%) simulation effort than the other two approaches. Also of note is the fact that the confidence interval surrounding the raw quantiles approach in conjunction with the known variance R&S procedure does not contain the PCS value of 0.95 set by the procedure, while the approach using Rinott's procedure

with raw quantiles does contain 0.95. This indicates the possibility that the known variance R&S procedure is more sensitive to non-normality in the comparison metric than the two-stage procedure which assumes unknown variance.

The confidence interval representing the macro-replications approach in Figure 2 is not dominated by any other points. Although the simulation effort for the macro-replications approach is considerably higher (more than twice as much) than the approach using the known variance of the means, the accuracy is also greater. Therefore, should a decision maker desire an extremely high realized PCS, he could choose to sacrifice simulation effort and use the macro-replications approach to achieve the increased level of accuracy. Overall, Figure 2 shows that the confidence intervals surrounding all the approaches, except the approach using quantiles directly with the known variance R&S procedure, contain the specified PCS value 0.95, but the grouped quantiles approach did so with considerably reduced simulation effort.

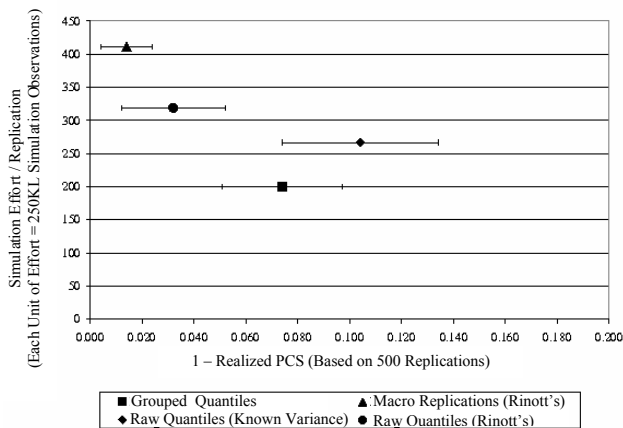


Figure 2: Simulation effort vs. realized probability of correct selection for all approaches.

4 CONCLUSIONS AND FUTURE WORK

In this paper, the variance of a group of quantile estimates is used to estimate the variance of a new variable surrounding the mean of a group of the same quantile estimates. The distribution surrounding the mean shares properties with the original set of quantile estimates, but also has the normal distribution required by the assumptions used to develop most R&S procedures. The variance of the variable surrounding the mean is used in conjunction with a R&S procedure to determine the required sample size from each system. The results in Section 3 highlight that the use of this approach allows quantiles to successfully be used as the comparison metric for ranking-and-selection procedures. The approach yielded high realized PCS values for the example given in this paper, while requiring significantly less simulation

effort than either the macro-replications or raw quantile approaches.

Additionally, the approach focuses on saving effort in the initial data collection phases of R&S procedures by reusing the original n_f simulation runs used to calculate the variance of the quantile estimates as observations for the R&S procedure. We believe these savings to be significant for systems such as semiconductor wafer fabs in which the variability of comparison metrics is not high, but the effort to obtain a single observation of that metric is very high. In such cases, the number of total observations needed from each system is not likely to be high, and the number of initial observations collected represents a large portion of the simulation effort. R&S procedures based on unknown variance require an initial sample size, $n_0 \geq 3$, and values of n_0 used in practice are often higher, making the initial effort required to get an estimate of the variance greater than the number of observations actually required to differentiate between systems. Alternatively, the approach used in this paper uses an initial set of n_f runs to obtain a precise estimate of the variance surrounding the performance measure of interest. In many cases, the number of simulation runs required to generate the n_0 observations, particularly if grouping of any kind is required to account for non-normality, is greater than the number of runs, n_f , required to precisely estimate the variance, resulting in an overall savings of simulation effort for these types of systems.

Also of note is the fact the R&S procedures with both known variance and unknown variances used in the empirical analysis in this paper are conservative. Rinott's procedure, for example, is not the most efficient R&S procedure in terms of required sample size and often allocates more samples than necessary to achieve the desired probability of correct selection (Chen and Kelton 2003). As a result, it is possible that other R&S procedures (with both known and unknown variances) would require less simulation effort than that illustrated in this paper.

Future work in this area will include the empirical evaluation of the procedures to other system configurations in which all systems are not exactly δ away from the best system. Additional evaluations in which the run length, m , is greater will also be performed. The run length of $m = 250k$ given in this paper was intentionally chosen to be fairly short so that the quantile estimates would be variable, increasing the simulation effort requirements. Future work will investigate how the realized PCS and simulation effort are comparatively affected by this run length parameter and by the configurations of the systems themselves.

ACKNOWLEDGMENTS

This research has been supported in part by grants 04-F1-1224 and 04-F1-1225 from the Factory Operations Research Center (FORCe) that is jointly funded by the Semiconductor

Research Corporation (SRC) and by International SEMATECH and in part by a doctoral fellowship sponsored by Intel Corporation and SRC.

REFERENCES

- Bechhofer, R.E., T.J. Santner, D.M. Goldsman. 1995. *Design and Analysis of Experiments for Statistical Selection, Screening, and Multiple Comparisons*. New York: John Wiley & Sons, Inc.
- Bekki, J., G. Mackulak, J. Fowler, and B.L. Nelson. 2006. Indirect cycle-time quantile estimation using the Cornish-Fisher Expansion. In review with *IIE Transactions*, available online at <www.fulton.asu.edu/~ie/research/workingpaper/wps.php> (ASU-IE-ORPS-2004-001).
- Chen, E.J. and W.D. Kelton. 2003. Inferences from indifference-zone selection procedures. In *Proceedings of the 2003 Winter Simulation Conference*, ed. S. Chick, P.J. Sánchez, D. Ferrin, and D.J. Morrice, 456–464.
- David, H.A. 1981. *Order Statistics, 2nd Edition*. New York: Wiley.
- Goldsman, D., B. Nelson, and B. Schmesier. 1991. Methods for selecting the best system. In *Proceedings of the 1991 Winter Simulation Conference*, ed. B. Nelson, W. Kelton, and G. Clark, 177–186.
- McNeill, J., G. Mackulak, and J. Fowler. 2003. Indirect estimation of cycle-time quantiles from discrete event simulation models using the Cornish-Fisher expansion. In *Proceedings of the 2003 Winter Simulation Conference*, ed. S. Chick, P.J. Sánchez, D. Ferrin, and D.J. Morrice, 1377–1382.
- Nelson, B.L. and F. Matejcek. 1995. Using common random numbers for indifference-zone selection and multiple comparisons in simulation. *Management Science* 41:1935–1945.
- Nelson, B.L., J. Swann, D. Goldsman, and W. Song. 2001. Simple procedures for selecting the best simulated system when the number of alternatives is large. *Operations Research* 49:950–963.
- Rinott, Y. 1978. On two-stage selection procedures and related probability inequalities. *Communications in Statistics* A7:799–811.
- Wilcox, R.R. 1984. A table for Rinott's selection procedure. *Journal of Quality Technology* 16:97–100.

AUTHOR BIOGRAPHIES

JENNIFER M. BEKKI is a Ph.D. candidate in the Industrial Engineering department at Arizona State University. Her research interests are in discrete-event simulation methodology and manufacturing applications in the semiconductor industry. Prior to beginning her Ph.D.

studies, she served as an intern in the Operational Decision Support Technologies group at Intel, and she is currently serving as an intern in the Assembly Test Technology and Development group. Jennifer is funded on the SRC/Intel Doctoral Fellowship.

JOHN W. FOWLER is a Professor of Industrial Engineering at Arizona State University (ASU) and is the Center Director for the Factory Operations Research Center that is jointly funded by International SEMATECH and the Semiconductor Research Corporation. His research interests include modeling, analysis, and control of semiconductor manufacturing systems. Dr. Fowler is a member of ASEE, IIE, INFORMS, POMS, and SCS. He is an Area Editor for SIMULATION: Transactions of the Society for Modeling and Simulation International and an Associate Editor of IEEE Transactions on Electronics Packaging Manufacturing. He is an IIE Fellow and is on the Winter Simulation Conference Board of Directors.

GERALD T. MACKULAK is an Associate Professor of Engineering in the Department of Industrial Engineering at Arizona State University. He is a graduate of Purdue University receiving his B.Sc., M.Sc., and Ph.D. degrees in the area of Industrial Engineering. His primary area of research is simulation applications within manufacturing with a special focus on automated material handling within semiconductor manufacturing.

BARRY L. NELSON is the Charles Deering McCormick Professor in the Department of Industrial Engineering & Management Sciences at Northwestern University. His research interests are the design and analysis of stochastic simulation experiments, particularly issues of multivariate input modeling, optimization via simulation and metamodeling. He is currently on the Board of Directors of the Winter Simulation Conference, and is a member of IIE (senior member), INFORMS (fellow), ACM, and ASA.