

CYCLE-TIME QUANTILE ESTIMATION IN MANUFACTURING SYSTEMS EMPLOYING DISPATCHING RULES

Jennifer E. McNeill
John. W. Fowler
Gerald T. Mackulak

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

Barry L. Nelson

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

ABSTRACT

The cycle-time distribution of manufacturing systems employing dispatching rules other than FIFO can be both highly skewed and have heavy tails. Previous cycle-time quantile estimation work has suggested that the Cornish-Fisher expansion can be used in conjunction with discrete-event simulation to provide cycle-time quantile estimates for a variety of systems operating under FIFO dispatching without requiring excess data storage. However, when the cycle-time distribution exhibits heavy skewness and kurtosis, the accuracy of quantile estimates obtained using the Cornish-Fisher expansion may degrade, sometimes severely. This paper demonstrates the degradation and motivates the need for a modification to the Cornish-Fisher expansion for estimating quantiles under non-FIFO dispatching rules. A solution approach combining a data transformation, the *maximum (minimum)-transformation*, with the Cornish-Fisher expansion is presented. Results show that this provides significant improvements in accuracy over using the Cornish-Fisher expansion alone while still retaining the advantage of requiring minimal data storage.

1 INTRODUCTION

An essential component of generating accurate lead time estimates, crucial for maintaining high levels of customer-service, is producing accurate estimates of cycle-time quantiles. Commonly, discrete-event simulation models of manufacturing systems are used to deliver estimates of the mean cycle-time. However, in many systems, quoting the mean cycle-time as the customer delivery-date results in a high percentage of late deliveries. As a result, it is preferable to use estimates of cycle-time quantiles. Quantiles provide the decision maker with greater detail about the cycle-time distribution than an estimate of the mean and

allow delivery-time quotes to be made with varying levels of confidence.

Unfortunately, obtaining cycle-time quantile estimates from discrete-event simulation models is difficult, often requiring excessive data storage of individual observations, advance knowledge of which quantiles estimates are required, or exhibiting estimation accuracy dependent on the cycle-time distribution itself. A quantile estimation technique that addresses all three of these issues simultaneously is currently unavailable. For example, Jain and Chlamtac (1984) give an algorithm for quantile estimation that does not require the storage of individual cycle-time observations. Instead, it requires advance knowledge of which quantiles are to be estimated, and if estimates of several quantiles of the same variable are required, the efficiency and accuracy of the algorithm degrades. Chen and Kelton (to appear), on the other hand, give a procedure that does not require prior knowledge of which quantiles are desired, but the approach has the drawback of requiring large sample sizes, especially for highly correlated systems. Finally, McNeill et al. (2005) provide a quantile estimation technique that has both low-data storage requirements and no condition for knowing ahead of time which quantiles estimates are desired. However, the accuracy of the technique is dependent on the distribution of the cycle-time values; quantile estimates from distributions close to the normal distribution have the greatest accuracy.

This paper builds on the work by McNeill et al. (2003, 2005) and presents an improvement to their approach, addressing the dependence that the accuracy of their approach has on the shape of the cycle-time distribution. Mathematical and empirical results are given which suggest that the proposed method has improved accuracy. Advantages and disadvantages of the approach are discussed, and directions for future work are suggested.

2 CORNISH-FISHER EXPANSION

McNeill et al. (2003) show that the Cornish-Fisher expansion can be used in conjunction with discrete-event simulation to estimate cycle-time quantiles in a variety of single-product manufacturing environments employing first-in-first-out (FIFO) dispatching rules at all workstations. The Cornish-Fisher expansion is an infinite series used to approximate normalized quantiles from any distribution, given a quantile from the standard normal distribution and the distribution's moments (Cornish and Fisher, 1937). Equation (1) gives a truncated version of the expansion, including only the first four terms. In this equation, z_α is a quantile from the standard normal distribution, γ_1 is the standardized central skewness, γ_2 is the standardized central excess kurtosis, σ is the standard deviation, μ is the mean, and Y^α is the quantile approximation. When sample moments are used in place of theoretical moments, Equation (1) becomes an estimator, and, consequently, it is important to have consistent, and ideally unbiased estimates of the first four sample moments.

$$Y^\alpha = \mu + \sigma x_\alpha, \text{ where } x_\alpha = z_\alpha + 1/6(z_\alpha^2 - 1)\gamma_1 + 1/24(z_\alpha^3 - 3z_\alpha)\gamma_2 - 1/36(2z_\alpha^3 - 5z_\alpha)\gamma_1^2 \quad (1)$$

To use the Cornish-Fisher expansion in conjunction with discrete-event simulation to obtain cycle-time quantile estimates, McNeill et al. (2005) suggest that estimates of the sample moments be calculated during the simulation run and then plugged into Equation (1) to obtain the quantile estimate. Minimal data storage, necessary for sample moment calculation, is required, and upon completion of simulation runs, any quantile of the distribution can be calculated without further simulation effort using Equation (1).

A shortcoming of the technique, however, presents itself when dispatching rules other than FIFO are utilized. In such cases, the cycle-time distribution may deviate drastically from the normal distribution. For example, the shortest-processing-time-first (SPT) dispatching rule causes both the excess kurtosis and skewness values at a given design point to explode. As jobs with short processing times quickly flow through each station, jobs with long processing times at a given machine are required to wait much longer than they would under a FIFO dispatching policy. As a result, an increased number of jobs have a very long cycle-time, resulting in a much heavier upper tail of the distribution. Consequently, the cycle-time distribution is much more heavily skewed toward its upper tail, and the excess kurtosis value is severely inflated over its FIFO counterpart.

Figures 1 and 2 illustrate the impact that dispatching rules can have on the cycle-time distribution of a simple system. Both figures were generated from simulations of an M/M/1 system at 90% utilization. The histograms and

the moment estimates were calculated based on the first 500,000 observations. Figure 1 represents the system under FIFO dispatching, while SPT was used to generate Figure 2. Additionally, common random numbers (CRN) were employed between the two systems. While a visual difference in the two histograms is present, the more dramatic impact of the implementation of SPT is on the moment estimates. While the mean under SPT is lower than under FIFO, as expected (Pinedo, 1995), all three of the higher moment estimates are dramatically higher.

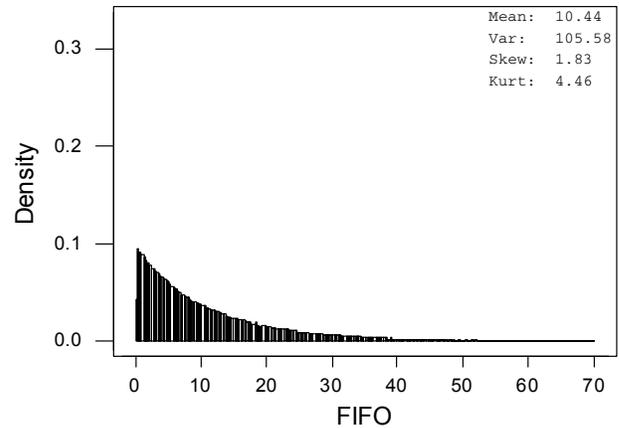


Figure 1: Cycle-time Histogram and Corresponding Moment-estimates for an M/M/1 System at 90% Utilization under a FIFO Dispatching Policy

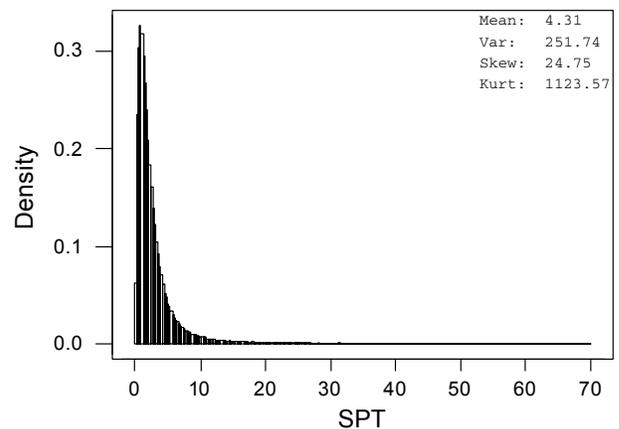


Figure 2: Cycle-time Histogram and Corresponding Moment-estimates for an M/M/1 System at 90% Utilization under SPT Dispatching

Furthermore, the higher moment estimates of the cycle-time distribution under SPT have a large impact on the accuracy that the first four terms of the Cornish-Fisher expansion have in estimating cycle-time quantiles from the distribution. This impact is illustrated in Figure 3. To generate the figure, 30 replications of 1,000,000 observations each were made of an M/M/1 system at 90% utilization. From each replication, estimates of the first four

moments were obtained, and the average of the 30 estimates for each moment were taken. These average moment estimates were then used in conjunction with Equation (1) to obtain cycle time quantile estimates for quantiles between 0.01 and 0.99. The x-axis of Figure 3 gives the quantile (0.01 to 0.99) being estimated while the y-axis gives the values of these quantiles estimated by using the average moment estimates with Equation (1). Figure 3 clearly illustrates accuracy problems that arise when using the Cornish-Fisher expansion in estimating quantiles from this distribution. Since the figure represents an estimate of the cumulative distribution function (cdf), it should be non-decreasing; instead, because of the extremely large moment values, it shows that the estimate of the 0.1 cycle-time quantile is significantly larger than the estimates of the 0.9 cycle-time quantile.

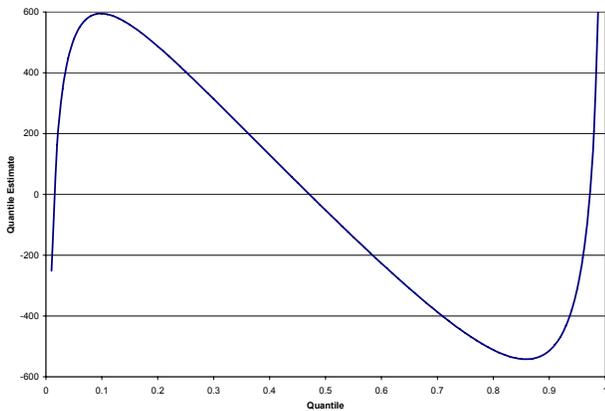


Figure 3: Quantile Estimates of an M/M/1 System at 90% Utilization under SPT Generated using the Cornish-Fisher Expansion

Experiments using the longest-processing-time-first (LPT) dispatching rule were also performed, and results similar to those shown in Figure 3 were found. It is also anticipated that other dispatching rules will yield comparable impacts on the cycle-time distribution and, therefore, on the accuracy of the Cornish-Fisher expansion in estimating cycle-time quantiles. Consequently, a modification to the approach suggested by McNeill et al. (2003, 2005) is clearly necessary when dispatching rules other than FIFO are in place. A combination of the maximum-transformation, proposed by Heidelberger and Lewis (1984), with the Cornish-Fisher approach is such a modification that shows promise in terms of an improvement in accuracy for non-normally distributed systems.

3 THE MAXIMUM/MINIMUM TRANSFORMATION

When applied to i.i.d. data, the maximum-transformation (or minimum transformation for quantiles less than 0.5) as proposed by Heidelberger and Lewis (1984) suggests that

for a sample of i.i.d. data, $x_1, x_2, x_3, \dots, x_n$, where n represents the number of independent observations from the x distribution, the p quantile ($p > 0.5$) can be estimated as the $q = p^v$ quantile of a sample of transformed data $y_1, y_2, y_3, y_4, \dots, y_m$, where each y_i value represents the maximum of a v number of x_i values and m represents the total number of derived values from the y distribution. For example, $y_1 = \max(x_1, x_2, \dots, x_v)$, $y_2 = \max(x_{v+1}, x_{v+2}, \dots, x_{v+v})$, and so on. The value of v depends on the quantile, p . Heidelberger and Lewis suggest using Equation (2) for determining v . Equation (2) guarantees that, regardless of the quantile estimated from the x_i samples, the quantile estimated from the y_i values will be close to the median. This is useful since the median of a sample distribution is generally considered to be easier to estimate than an extreme quantile from either tail of the distribution.

$$v = \left\lceil \frac{\ln(0.5)}{\ln(p)} \right\rceil \tag{2}$$

Notice in Equation (2) that as the quantile from the x_i sample distribution approaches 1, the value of v approaches infinity. Moreover, as p tends towards 1/2, v tends towards 1. When $v=1$, there is no difference in estimating quantiles directly from the x_i values or the y_i values, since in this case each y_i value is simply the “maximum” of a single x_i value. For reference, Table 1 gives five p values, and their resulting v values. The v values are not dependent on the underlying population of the x_i samples. Table 5 shows, for example, that regardless of the true distribution from which the x_i samples were drawn, to estimate the 0.99 quantile, 68 x_i values are required to generate a single y_i value.

Table 1: Number of x_i Values (v) used to Generate Each y_i Value Via the Maximum Operator for Various Quantiles (p) of the x_i Distribution

p	v
0.7	1
0.8	3
0.9	6
0.95	13
0.99	68
0.999	392

4 MAXIMUM/MINIMUM TRANSFORMATION WITH THE CORNISH-FISHER EXPANSION

4.1 Exponential Distribution

To understand the potential impact of combining the maximum (or minimum) transformation with the Cornish-Fisher expansion for estimating cycle-time quantiles, the theoretical

moments of the y_i distribution were calculated when the x_i distribution is exponentially distributed with a λ value of 1. Since quantile estimates obtained from the exponential distribution using the first four terms of the Cornish-Fisher expansion alone are known to degrade for quantiles lower than 0.5, getting worse as the estimated quantiles get closer to the extreme lower tail (McNeill et al., 2005), the minimum transformation was used, and quantiles from the lower tail of the distribution were estimated. Using the minimum transformation, the p quantile of the x_i values ($p < 0.5$) is estimated as the $q=1-(1-p)^v$ quantile of the y_i values. Additionally, when using the minimum transformation, the formula used to determine v is slightly different than for the maximum-transformation and is given in Equation (3).

$$v = \left\lfloor \frac{\ln(1 - 0.5)}{\ln(1 - p)} \right\rfloor \quad (3)$$

To use the Cornish-Fisher expansion with the minimum transformation, calculation of the moments of the y_i values was required. The derivation of the CDF and the probability density function (PDF) are given below, followed by the resulting raw moment values, μ'_k , shown in Equation (4). In Equation (4), k indicates the k th moment.

Assume $x_i \sim \text{Exp}(\lambda)$. Then

$$P(Y \leq y) = P(x_1 \leq y \text{ or } x_2 \leq y \text{ or } \dots \text{ or } x_v \leq y)$$

$$P(Y \leq y) = P(\text{at least } 1 \text{ } x_i \text{ value } \leq y \text{ for } 1 \leq i \leq v)$$

$$P(Y \leq y) = 1 - P(\text{zero } x_i \text{ values } \leq y \text{ for } 1 \leq i \leq v)$$

$$P(\text{zero } x_i \text{ values } \leq y) = (1 - e^{-\lambda y})^v$$

$$\text{CDF: } P(Y \leq y) = 1 - (1 - e^{-\lambda y})^v$$

$$\text{PDF: } v\lambda(e^{-\lambda y})^v$$

$$\mu'_k = \int_0^\infty y^k v(e^{-\lambda y})^v dy = k! \frac{\lambda^k}{v^k} \quad (4)$$

Using the moment values calculated from Equation (4) with the first four terms of the Cornish-Fisher expansion (Equation (1)), estimates of quantiles between 0.05 and 0.5 of the x_i distribution were made. Estimated quantiles were compared with the theoretical values, and the percentage difference between the two values is reported on the y-axis of Figure 4. Since the moments are known, rather than estimated from data, any error is attributable to the Cornish-Fisher expansion. The same quantile estimates were also made using the Cornish-Fisher expansion without first transforming the data, and the percentage difference between these values and the true quantiles from the exponential distribution is also reported in Figure 4. When the estimated quantile is above 0.3, the v value is 1, making the two estimation techniques equivalent. In the extreme lower tail, Figure 4 shows that the minimum transformation results in a significant gain in accuracy.

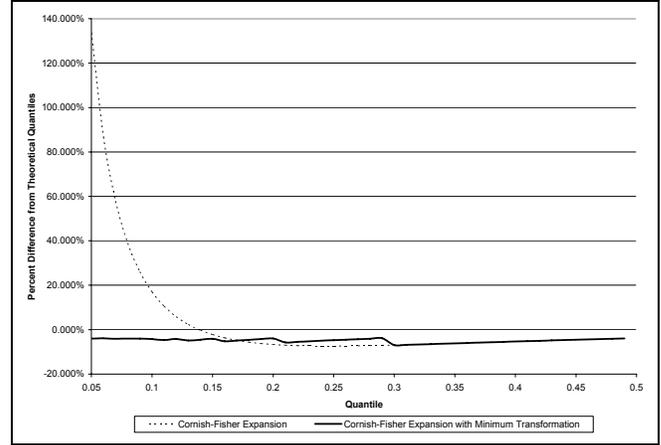


Figure 4: Impact on Quantile Estimation Accuracy of Combining the Minimum-transformation with the Cornish-Fisher Expansion when the x_i Distribution is Exponentially Distributed with $\lambda=1$

4.2 Lognormal Distribution

To further illustrate potential gains in accuracy when estimating quantiles from distributions with high moment values, experimentation was done on i.i.d. data generated from a lognormal distribution. With a C++ program, the 0.8, 0.9, 0.99, and 0.99 quantiles were estimated from a lognormal (0.75,1) distribution with a shape parameter of 0.75 and a scale parameter of 1 using the described maximum-transformation. This distribution was selected since its moment values are high ($\gamma_1 = 3.26$ and $\gamma_2 = 23.54$), causing the Cornish-Fisher expansion applied directly (i.e., no transformation) to have poor accuracy in estimating its quantiles.

Random, i.i.d. samples were drawn from the distribution until there were 1000 y_i values; the total number of x_i samples, N , to estimate a given quantile is $N = v * 1000$. For instance, 6000 (=1000*6) random samples from each distribution were drawn to estimate the 0.9 quantile, while only 3000 were required to estimate the 0.8 quantile. To estimate each quantile of each distribution, 5 independent samples of N x_i values were drawn. For each of the five independent samples of N values, estimates of the first four moments of the y_i distribution were calculated. Using these sample moment estimates in Equation (1), the p quantile from the x_i sample distribution was estimated as the $q=p^v$ quantile of the y_i distribution. Since the random samples were drawn from populations of known distributions, the theoretical quantiles from those distributions were also available. Therefore, it was possible to calculate the relative percentage difference between the estimated and true quantile. This value was calculated for each of the five independent replications performed for each quantile (0.8, 0.9, 0.99, and 0.999). The average of these five percentage differences was recorded as the average percentage

error in estimating a given quantile from a given distribution using the maximum (minimum) - transformation in conjunction with the Cornish-Fisher expansion on i.i.d. data.

In addition to knowing the theoretical quantiles from the test distributions, the theoretical moments were also obtainable. Using these values, the “best case” performance of the Cornish-Fisher expansion *without* the use of the maximum-transformation was calculated by approximating each of the same quantiles using the theoretical moments of each distribution in Equation (1). The quantile approximations obtained with the true moments were then compared with the true quantiles from the same distributions, and the relative percent difference between the two was recorded. These results were considered “best case” as they reflect the level of accuracy that the Cornish-Fisher expansion could achieve in the absence of any sampling-induced error.

The “best case” results without the use of the maximum-transformation were then compared to the average percentage errors obtained using the maximum-transformation. The direct difference between the two was calculated and is reported in Table 2. A positive value in the “average % accuracy gain” column should be interpreted as indicating that, on the average, using the maximum-transformation resulted in an accuracy improvement over the best case results without the transformation. For instance, in estimating the 0.8 quantile, the use of the maximum-transformation in conjunction with the Cornish-Fisher expansion resulted in a 35% improvement in accuracy over using the theoretical moments directly in the Cornish-Fisher expansion without first using the transformation. A negative sign in this same column indicates the reverse situation; in these cases, on the average, the maximum transformation resulted in a decrease in estimation accuracy. To give a feel for the range, the maximum and minimum percent accuracy gains across the five simulated replications are also reported in Table 2.

Table 2: Accuracy Gain in Estimating Quantiles from the Lognormal (0.75,1) Distribution using the Maximum-transformation in Conjunction with the Cornish-Fisher Expansion over using the Cornish-Fisher Expansion Alone.

Quantile	Average % Accuracy Gain	Minimum % Accuracy Gain	Maximum % Accuracy Gain
0.8	35%	21%	47%
0.9	5%	-29%	21%
0.99	42%	38%	44%
0.999	78%	77%	79%

Table 2 clearly shows a marked increase in accuracy after using the maximum-transformation for the lognormal (0.75,1) distribution. Additionally, Table 2 shows that, with the exception of the 0.9 quantile, as the estimated quantile gets closer to 1.0, the improvement in accuracy

increases. This is likely explained by the fact that as the quantile being estimated gets closer to 1.0, the value of ν also increases, making the impact of the maximum operator greater.

5 DISCUSSION AND CONCLUSIONS

The results highlight the potential benefits in terms of accuracy of using the maximum (minimum) -transformation in conjunction with the Cornish-Fisher expansion to estimate cycle-time quantiles. The largest advantage over using the Cornish-Fisher expansion without the data transformation is clearly the accuracy improvement, particularly for extreme quantiles and for distributions with high skewness and kurtosis values. These results are particularly of note since without the maximum-transformation, the first four terms of the Cornish-Fisher expansion do a poor job at estimating quantiles from these same distributions. Furthermore, assuming that the high skewness and excess kurtosis values contribute to the poor estimation performance of the expansion without the transformation, the expectation is that the same type of accuracy improvements will apply to cycle-time distributions from manufacturing settings in which a dispatching rule other than FIFO (i.e., SPT) is employed. Accuracy improvement through the use of the maximum-transformation operator would be extremely useful in such a system.

With the clear advantages of the approach in terms of accuracy comes the disadvantage of estimating multiple quantiles from the same set of simulation runs. Since the ν value dictates the q quantile estimated from the y_i sample distribution, and the ν value depends directly on the p quantile desired from the original x_i distribution, obtaining estimates of different p quantiles requires different ν values (assuming the intent is to estimate the $q=0.5$ quantile from the y_i distribution). As a result, *a priori* knowledge of which quantile estimates are desired would be required so that separate sets of y_i distributions can be maintained, each set of which is based on a different ν value. Alternatively, the ν value for the highest quantile could be used for all quantiles, requiring only moment estimates from single y_i distribution to be maintained, but requiring estimates of different q values from that y_i distribution. Provided that the Cornish-Fisher expansion estimates all quantiles from a distribution with equal accuracy, this would be a reasonable solution. However, it is known that the accuracy of the Cornish-Fisher expansion depends greatly on the quantile being estimated for many systems (McNeill et al., 2003), and, as a result, estimates of q for varying quantiles would also have varying amounts of error induced by the expansion itself. (Of course, this variable estimation accuracy could also be exploited by selecting ν so that q is always a quantile estimated by the Cornish-Fisher expansion with high accuracy).

Finally, Figure 5 presents a taxonomy of four quantile estimation techniques (order statistics, maximum (mini-

mum)-transformation, and the Cornish-Fisher expansion with and without the use of the maximum-transformation) in terms of their data storage requirements and ability to estimate multiple quantiles. Although not previously discussed in this paper, order statistics are included in the comparison as they represent the most traditional quantile estimation technique. To generate an estimate of a cycle-time quantile using order statistics all data points are stored, sorted, and the appropriate quantile is selected from the sorted values.

Figure 5 shows that both order statistics and the maximum-transformation approach require storage of individual observations from the sample distribution. Note that only order statistics require storage of all observations. The maximum-transformation requires a storage size equivalent to $1/v$ the size required using only order statistics, but the individual y_i observations must still be saved and sorted to implement the procedure. Both versions of the Cornish-Fisher expansion, on the other hand, provide the advantage of requiring very little data storage. Additionally, if multiple quantile estimates are desired, both the maximum-transformation alone and the maximum-transformation with the Cornish-Fisher expansion require advance knowledge of which quantiles are to be estimated. Conversely, order statistics and the Cornish-Fisher expansion without the maximum-transformation do not require any advance knowledge about which quantiles are desired. Using order statistics, however, some post-processing in the form of parsing the sorted list of observations may be required if desired quantiles are not specified in advance. The Cornish-Fisher expansion without the maximum-transformation also provides the great benefit of being able to generate any quantile estimate (or a discrete estimate of the CDF of the distribution) without any previous knowledge of which quantiles are desired. Unfortunately, when the maximum-transformation is used in conjunction with the Cornish-Fisher expansion, Figure 5 shows that this benefit is lost, assuming the method is implemented as it has been described by this study. If, however, a reasonable solution to the issue of obtaining multiple quantile estimates without prior knowledge of which quantiles should be estimated can be found, the improved accuracy obtained by adding the maximum/minimum-transformation to the Cornish-Fisher expansion makes it an extremely attractive cycle-time quantile estimation technique.

		Data Storage Requirements	
		Store actual samples from distribution	Store distributional parameters
Estimating multiple quantiles	Quantiles known in advance	Maximum Transformation	Cornish-Fisher Expansion WITH maximum (minimum) transformation
	Advance knowledge Not required	Order Statistics	Cornish-Fisher Expansion

Figure 5: Taxonomy of Quantile Estimation Techniques

6 FUTURE WORK

Results presented in this paper showed that when the skewness and kurtosis estimates of a distribution differ greatly from those of the normal distribution, the ability of the Cornish-Fisher expansion alone to accurately estimate quantiles from that distribution degrades substantially. Cycle-time distributions of systems under some dispatching rules clearly fall under this category. Future work should include such results for a variety of dispatching rules, and these quantile estimates should be compared to estimates obtained from very long simulation runs. Of note is the fact that the data transformation in this case would be applied to simulation-generated data, which is not i.i.d. If, however, the cycle-time distribution generated from a simulation model is assumed to be made of stationary, dependent data, the maximum (or minimum) data transformation can still be performed, and Heidelberger and Lewis (1984) present a modification to the approach discussed previously for i.i.d. data.

Additional future work should also include an investigation of ways to harness the accuracy improvement that the combination of the maximum-transformation in conjunction with the Cornish-Fisher expansion provides without losing the benefit of not requiring which quantile estimates are desired to be known ahead of time. It is possible that the data transformation could be used to obtain estimates of specific quantiles from the cycle-time distribution, and these estimates could then be used to parameterize the Cornish-Fisher expansion so that it is customized for each distribution and able to better estimate quantiles from even extremely non-normal distributions.

ACKNOWLEDGMENTS

This research has been supported in part by grant P011335 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

- Chen, E.J. and Kelton, W.D. (to appear) Quantile and Tolerance-Interval Estimation in simulation. To appear: *European Journal of Operational Research*.
- Cornish, E.A. and R.A. Fisher. 1937. Moments and Cumulants in the Specification of Distributions. *Revue de l'Institut International de Statistique*, 5: 307-320
- Heidelberger, P. and Lewis, P. 1984. Quantile Estimation in Dependent Sequences. *Operations Research*. 32(1): 185-209.
- Jain, R. and Chlamtac, I. 1985. The P² Algorithm for Dynamic Calculation of Quantiles and Histograms Without Storing observations. *Communications of the ACM*, 28(10): 1076-1085.
- McNeill, J.E., Mackulak, G.T., and Fowler, J.W. 2003. Indirect Estimation of Cycle-time Quantile From Discrete-event Simulation Models Using the Cornish-Fisher Expansion, *Proceedings of the 2003 Winter Simulation Conference*, S. Chick, P. J. Sánchez, D. Ferrin, and D. J. Morrice, eds., 1377 – 1382.
- McNeill, J.E., Mackulak, G.T., Fowler, J.W., and Nelson, B.L. 2005. Indirect Cycle-Time Quantile Estimation Using the Cornish-Fisher Expansion, *Working Paper, Arizona State University*. Available online at: <http://www.eas.asu.edu/ie/workingpaper/wps.php>.
- Pinedo, M. (1995) *Scheduling Theory, Algorithms, and Systems*. Prentice Hall: New Jersey.

AUTHOR BIOGRAPHIES

JENNIFER E. MCNEILL 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 funded on an 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 James N. and Margie M. Krebs 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 Chair of the Board of Directors of the Winter Simulation Conference.