

DYNAMIC OUTPUT ANALYSIS FOR SIMULATIONS OF MANUFACTURING ENVIRONMENTS

Peter Mullarkey

Maxager Technology, Inc.
2173 E. Francisco Blvd., Suite C
San Rafael, CA 94901, U.S.A.

Srinagesh Gavirneni

Maxager Technology, Inc.
8804 Rattlebush Cove
Austin, TX 78750, U.S.A.

Douglas J. Morrice

Graduate School Of Business
University of Texas at Austin
Austin, TX 78712, U.S.A.

ABSTRACT

We describe the design and implementation of a generic, real-time, in-line output analysis procedure for controlling simulations of discrete manufacturing environments. We implemented this capability in the commercial simulation software Extend®. The main issues we faced were (1) Specifying the products to evaluate, (2) Determining the batch sizes for output analysis, and (3) Defining the stopping conditions based on the confidence intervals. We implemented a significance test for correlation and used this test to dynamically adjust the batch sizes used in confidence interval estimation done using batch means. When the stability conditions have been met, the simulation prompts the user to consider stopping the simulation. On the other hand, if at the end of the run length selected by the user, the statistical conditions were not satisfied, the tool notifies the user of that fact. This capability enabled us to significantly reduce the simulation run lengths, and ensures, with little additional computational effort, that the results were reliable. We used this tool to control simulations of electronics, steel, automotive, and metal processing industries. In general, using this tool we realized a reduction of more than 40% in the time required for simulation.

1 INTRODUCTION

Simulation is a popular methodology for analyzing complex manufacturing environments. It is especially useful in the presence of many processes, multiple products, significant setup times, complicated routings, and randomness in capacities, processing times, lead times, and

yield. Often, simulation is the only available approach for answering questions such as:

1. What is the effect of increasing the capacity at my bottleneck process?
2. How does a change in product mix effect my revenues?
3. Is it beneficial to invest in improving the yield at a process stage?

These are issues that we encounter in our work at Maxager. Maxager is the leading business-to-business software solution that enables manufacturers to sell their capacity most profitably in both traditional offline markets and emerging online eMarkets.

When attempting to answer these questions using simulation, one must take care to ensure that the results are statistically valid. Thus, there has been significant effort in developing sophisticated output analysis procedures and they are available in most commercial simulation packages such as Arena, ProModel, and Witness. Our work makes the following contributions:

1. Provides an integrated output analysis procedure for simulations in Extend. The design and implementation of this capability is generic enough to be easily implemented in any other simulation software.
2. Our procedure allows the user to select a performance measure (e.g., cycle time at a machine, cycle time through the factory, etc.) and specify independent stopping conditions based on the observations of a selected subset of products.

3. Many times, the analyst may be an expert in the manufacturing process, but may not be well versed in statistical analysis. Our procedure has the distinct advantage that the output analysis procedure requires little interaction from the analyst while maintaining rigor in statistical analysis.

In order to provide this functionality, we developed a procedure that is well integrated into the simulation model, requires little interaction from the analyst, dynamically changes its parameters as the simulation proceeds, prompts the analyst when the statistical conditions have been met, or extends the run length if it is too short. We implemented this tool in Extend®, a high quality commercially available, extensible, simulation software tool from ImagineThat Inc. However, the design and the algorithms are generic enough, so that they can be implemented in any simulation package.

We used this procedure to control simulations of manufacturing processes in the steel, electronics, rubber, and metal processing industries. Having this capability enabled us to ensure that our analysts were using the results from long-enough simulations to make decisions. In addition, we realized significant reductions (more than 40%) in the computational time required for simulation.

The rest of this paper is organized as follows. In Section 2, we detail our analysis procedure. We present a case study from a steel manufacturer in Section 3, and conclude in Section 4.

2 ANALYSIS PROCEDURE

This procedure contains three main features. They are:

1. **Product Importance Table:** Lets the analyst specify a subset of products for which the simulation run should be long enough to provide reliable statistics.
2. **Significance Test for Correlation:** For any given batch size, we use a Student’s *t*-test to check the significance of correlation between the batches, use the Fishman’s test (Fishman 1978) to check for independence between batch means, and increase the batch size if necessary.
3. **Confidence Interval Estimation:** For each important product, we compute the 100(1- α)% confidence interval and ensure that it meets the tightness specified by the analyst.

We explain each of these features in the following subsections.

2.1 Product Importance Table

In many situations, even if the manufacturing environment contains a large number of products, the analyst may be interested in the behavior of only a subset of products. Then, it is only necessary to ensure that the simulation is long enough to provide sufficient data for each of the products in the subset. We provide this ability by using a product importance table. In this table, for each product, the analyst enters a ‘1’ in the importance column if the product is important or enters a ‘0’ if it is not. Since the computational impact is low, we perform all the calculations for every product, but change the batch size only when necessitated by one of the important products.

2.2 Significance Test for Correlation

Since we are using the ‘Batch Means’ method (see Law and Kelton, pp. 554-555) for computing the confidence interval, we must make sure that the batch size is large enough to make successive batches independent. To check the significance of the correlation, we use a *t*-test. Suppose we are using a batch size of *B*, then:

1. Compute the correlation coefficient, *r*, between observations *i* and *B+i* for all the *N* available observations.
2. Compute the *t*-statistic using the formula:

$$t = \frac{r\sqrt{N-B-2}}{\sqrt{1-r^2}}$$
3. Test for significance by comparing this *t*-statistic to $t_{N-B-2, \frac{\alpha}{2}}$ for a test with confidence level 100(1- α).

We perform this test for all the important products and if the test fails for any of them, we increase the batch size for all of them. We start with an initial batch size of 10 for all the important products and increase the batch size by 10 every time the correlation coefficient is significantly different from zero. Once we have established that the successive batches are uncorrelated, we test for independence among batch means using Fishman’s test as detailed below. Assume that we have *n* batches of *B* observations in each batch. Let X_{ij} be the *j*th observation in the *i*th batch.

1. We compute the mean for batch *i*,

$$\bar{X}_i = \frac{1}{B} \sum_{j=B(i-1)+1}^{j=Bi} X_{ij} \text{ and the overall mean,}$$

$$\bar{\bar{X}} = \frac{1}{n} \sum_{i=1}^n \bar{X}_i .$$

2. Compute the correlation, r^b , between successive batch means using the formula:

$$r^b = \frac{\sum_{i=1}^{n-1} (\bar{X}_i - \bar{X})(\bar{X}_{i+1} - \bar{X})}{2 \sum_{i=1}^n (\bar{X}_i - \bar{X})^2}$$

3. Compute the test statistic, C_n , using the formula:

$$C_n = r^b + \frac{(\bar{X}_1 - \bar{X})^2 + (\bar{X}_n - \bar{X})^2}{2 \sum_{i=1}^n (\bar{X}_i - \bar{X})^2}$$

4. We conclude that the correlation is significantly different from zero if $C_n \sqrt{\frac{n^2-1}{n-2}}$ is greater than $Z_{\frac{\alpha}{2}}$, where $Z_{\frac{\alpha}{2}}$ is the standard normal variable with $\frac{\alpha}{2}$ in the right tail.

If the correlation is significantly different from zero, we increase the batch size by 10 for all the products and run the simulation longer to collect the necessary observations, and then retest.

2.3 Confidence Interval Estimation

We compute the confidence interval for the mean using the batch means method.

1. Compute the standard deviation of the batch means,

$$S_{\bar{X}} = \sqrt{\frac{1}{n-1} \left(\bar{X}_i - \bar{X} \right)^2}$$

2. Compute the 100(1- α)% confidence interval as:

$$\bar{X} \pm t_{n-1, \frac{\alpha}{2}} \frac{S_{\bar{X}}}{\sqrt{n}}$$

We use 10 batches to compute the confidence interval and hold this number constant. This is to ensure that we use fewer batches and larger size as recommended in Schmeiser (1982). We start with a batch size of 10 and increase it gradually until the independence between batches has been demonstrated.

Under very general conditions, as the batch size increases, the confidence interval for the mean gets tighter and we allow the analyst to specify a parameter, γ , such that the half-width of the confidence interval is less than $\gamma \bar{X}$ for the important products. After the simulation has run long enough to ensure independence between

successive batches, if the confidence intervals are not tight enough, we determine the batch size required to achieve this tightness in the confidence interval by setting it higher

than $\left(t_{n-1, \frac{\alpha}{2}} \frac{S_X}{\sqrt{n\gamma \bar{X}}} \right)^2$ where S_X is the standard deviation

of the individual observations. Note that we are using the approximation, $S_{\bar{X}} = \frac{S_X}{\sqrt{B}}$, in order to estimate the batch size. In most simulation runs, with positively correlated data, this approximation is likely to underestimate the actual batch size (see Law and Kelton, page 555) required to achieve the pre-specified precision. Thus we use this estimate as a starting point and increase the batch size as needed to achieve the required confidence interval tightness.

The complete flowchart for our procedure is given in Figure 1. We start with an initial batch size of 10 and run the simulation until we have observed ten batches for all the important products. As we collect the final sample for the tenth batch of an important product, we test for independence between batches and the batch means and increase the batch size if necessary. Once the independence has been observed for all the important products, we compute the confidence intervals and check if they satisfy the tightness criterion specified by the analyst. If they do, we (optionally) prompt the user to consider stopping the simulation run since we have sufficient data to draw reliable conclusions.

If the confidence intervals are not tight enough, we compute the batch size required to achieve the tightness and continue to run the simulation until we have ten batches of the new size for all the important products. If at any time, the simulation is about to stop since the run length selected by the analyst has been reached, we prompt the user to consider allowing the system to automatically increase the run length since the statistical requirements have not been met.

The user interface that we developed in Extend® for this analysis is given in Figure 2. The analyst can specify the performance measure on which we should do this analysis. Currently we have two choices for the performance measure, (1) cycle time at a pre-specified bottleneck machine, or (2) cycle time though the whole factory. We analyze the results on a product-specific basis, but it is easy to add a feature that performs the analysis on machine-based results, such as utilization, downtime etc. There is also the ability to clear statistics after an analyst specified warm-up period, in order to avoid bias from an initial transient.

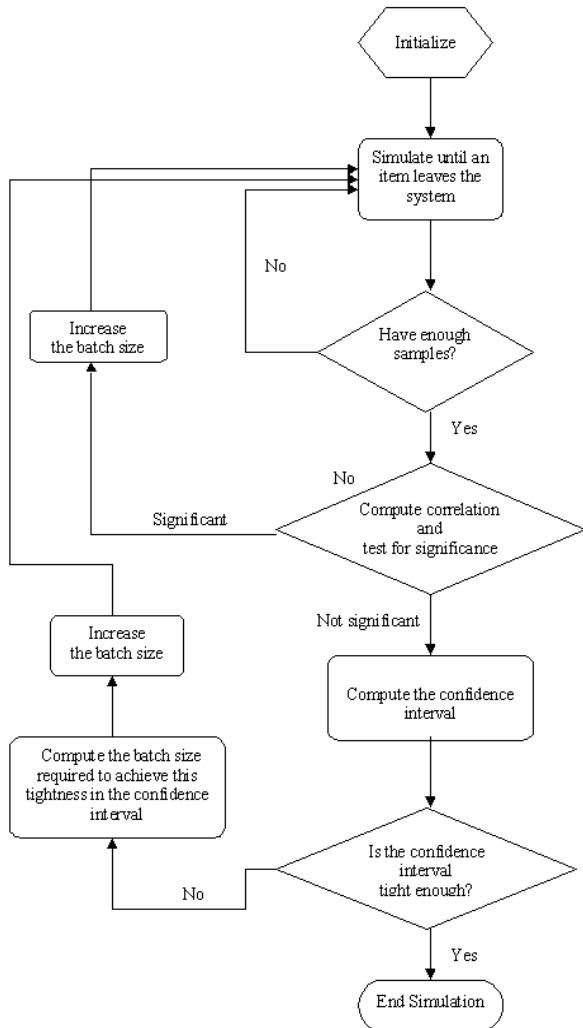


Figure 1: The Flow Chart For The Output Analyzer

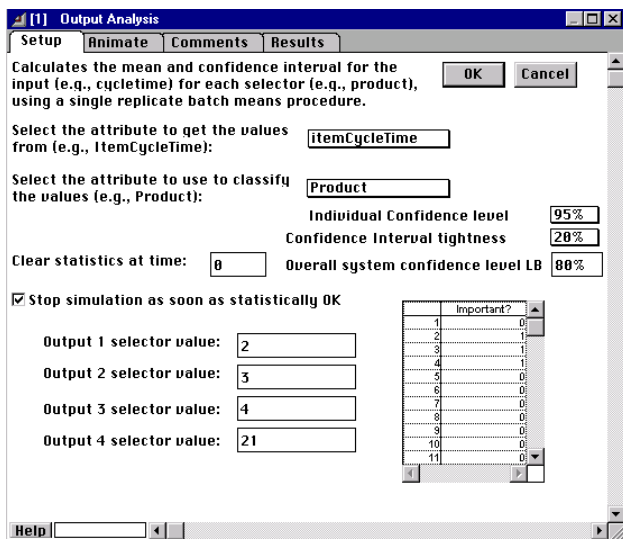


Figure 2: The User Interface for Specifying Stopping Conditions

The analyst can choose one of three possible values (5, 10, or 20%) for confidence interval precision. There are three choices for confidence level (95, 97.5, or 99%) the analyst can choose from, for the interval estimation. It is worth noting that this confidence level is for each product separately. The resulting lower bound for confidence level for the system (based on the number of important products), computed using the Bonferroni inequality (Law and Kelton, pp. 568-569), is displayed right below it. The analyst must use this number to judge his/her decisions if they require a certain performance for all the products simultaneously.

At the bottom right corner, the analyst specifies the important products and at the bottom left corner, he/she can choose a subset of these important products that may need some review. The results of this selected subset of products will be illustrated in histogram format at the end of the simulation.

In the next section we detail a case study of a steel manufacturing facility and explain how we used this procedure to control its simulation and analysis.

3 CASE STUDY: A STEEL PLANT

We used this capability to control simulations of manufacturing environments in electronics, steel, metal processing, sheet metal fabrication, and rubber industries. In all the cases, the tool was instrumental in ensuring that the time for simulation and analysis was as short as possible and that the analysts were using reliable data to make important decisions about product mix, capacity planning, and pricing. In this section, we explain in detail our usage of this tool in controlling the simulation of a large steel plant. In order to not reveal any company sensitive information, we have kept the descriptions fuzzy and left out many details. The intention of this case study is to give the reader an idea of how this capability would be used in a general setting and describe the type of benefits one can expect. First we describe the manufacturing setting in terms of the products, processes, and the market conditions. Then we briefly describe the needs that motivated the use of simulation. Later we present the results from using this tool and compare them to the results that would have been observed if this capability was absent.

3.1 Products

This steel plant manufactured and sold 31 different products from various stages of its manufacturing process. Out of these 31 products, four of them accounted for 80% of the total revenue at the plant. Thus whenever strategic decisions were being made, it was necessary to understand in detail the effect those decisions will have on these four products. While the impact on the other twenty-seven products is also important, it is usually a secondary concern.

3.2 Processes

This steel plant consists of six major process steps: Hot Mill, Pickling, Cold Rolling, Annealing, Tempering, and Galvanizing. At each step, there are varying degrees of setup, downtime and scrap. The steel starts as slabs, which go from the Caster to the Hot Rolling Mill. Some of the output of Hot Mill is sold directly as Hot Band. The rest of the coils go into Pickling. After Pickling, coils are routed to either Galvanize or Cold Rolling. After Galvanize, the product is sold as Galvanized Coils To Trade. The coils routed through Cold Rolling go to either Annealing and the Temper Mill, or Galvanize. The products sold after Temper Mill are Cold Rolled Coils To Trade. As can be seen, there is quite a bit of product diversity and routing.

3.3 Market Conditions

Recently, due to the drop in steel prices and stiff competition from smaller, more responsive mini-mills, major worldwide players have seen a significant drop in revenues and profits. Thus it was necessary to analyze their manufacturing process in detail to understand the profitability of the various products it markets. Once the profitable products have been identified, then the product mix must be altered to make sure that plant produces more of the high profit products and less of the low profit products. In addition, it may also be possible to change the pricing of these products and realize more profits. To understand these issues and the effect of these decisions, we built a simulation model (see Figure 3 below) for a large steel company. The simulation models the

manufacturing process and the uncertainties that are present in it.

3.4 Results

In this section we report on the results we observed while using this output analyzer in the simulation of the steel plant described above. We will discuss the impact of the three main features (product importance table, correlation significance test, and confidence interval tightness) of the Output Analysis block.

3.4.1 Effect of Product Importance Table

If the analyst did not have this importance table, he/she may try to collect 2000 data points (industry standard of 10 batches of 200 observations as recommended in Schmeiser (1982)) for each of the products. With the modeled product mix that would require a simulation run length of over 25 years needing a computational time of about six hours. If the analyst were to specify a subset of important products, only for which we require 2000 observations, then the simulation run length as well as the simulation time will be much shorter. For example, when the analyst decided to concentrate only on the four high volume products (2, 3, 4, and 21) we observed that we needed only 100 days of simulation run length requiring a computation time of less than 5 minutes. Thus the presence of the product importance table allows the analyst to focus on the products important to their study and drastically reduce the time required for simulation.

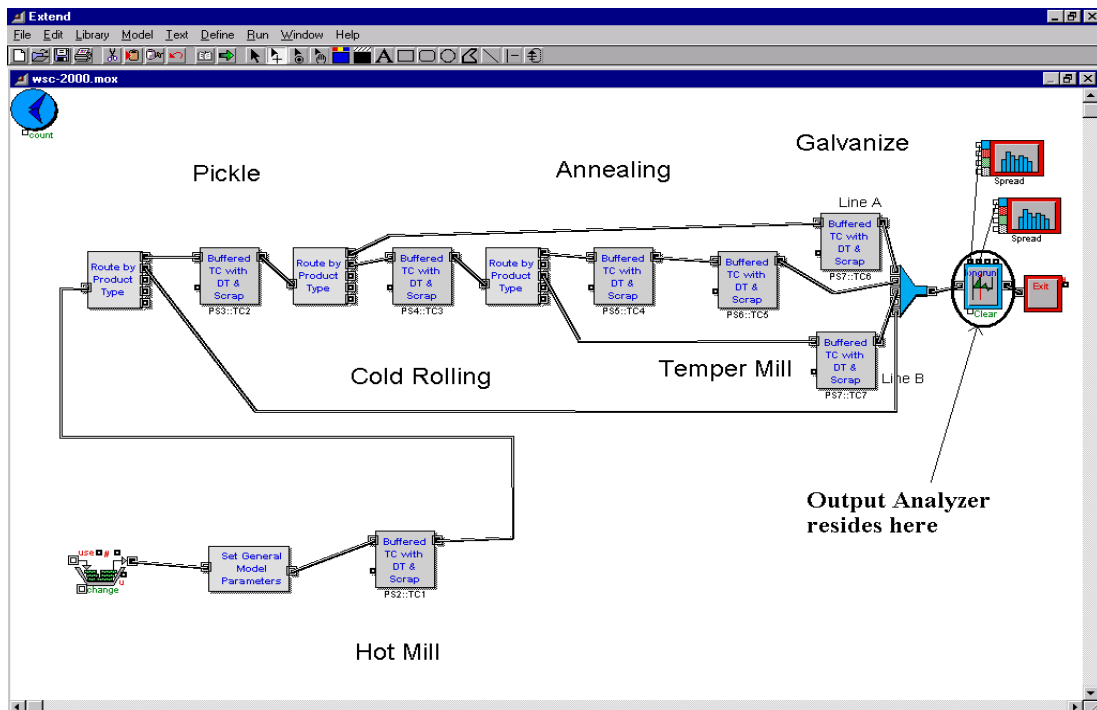


Figure 3: The Simulation Model of the Steel Plant

3.4.2 Effect of Correlation Significance Test

The industry standard guidelines were defined in a way to ensure statistical stability given the usual amount of intra-batch correlation. Often, these standards are very conservative and the correlation significance test helps us, if possible, use a smaller batch size. In this case, using the correlation significance test, we were able to dynamically change the batch size starting at 10 and increasing each time by 10. We observed that the independence between the batches was achieved at a batch size of 30 thus necessitating only 300 observations for each important product. This resulted in the simulation run length dropping from 100 days to 15 days, further saving computational effort.

3.4.3 Effect of Confidence Interval Tightness

When we simulated the system for 15 days and computed the confidence interval for the mean, we observed that the ratio of the half-width to the mean ranged from 2.3 to 14.2% for these four important products. Thus, if the analyst was looking for confidence interval tightness better than 20%, they would have been satisfied with the current results. However, if the analyst wanted a tightness of 5%, three of the products would have violated it, necessitating that the analyst run the simulation (of a longer length) at

least one more time. However, with the ability to specify the required tightness, the analyst, using our tool, can achieve these results using just one simulation run. If the analyst specified the tightness at 5%, we found that with the existing variance in the cycletime, the tool would automatically compute that it would need a batch size of 150. The simulation would then be run until the necessary observations have been recorded. The resulting confidence intervals had tightness ranging from 1.7% to 4.7% satisfying the restriction specified by the analyst. Thus, with this capability we have removed the need for repeat simulations.

If necessary, the analyst can take a closer look at the observations for a specific product using a histogram format. For example, the histogram representing the data for product 21 is given in Figure 4 below. We observed that this product was behaving as expected and it is not necessary to perform additional analysis. However, if there were anomalies, the histogram would have helped us discover them.

We have found that by encapsulating the decisions about batch size, numbers of batches, and most significantly run control, into the internals of the block, while having a few controls available in the user interface, our users have adapted to it quite easily.

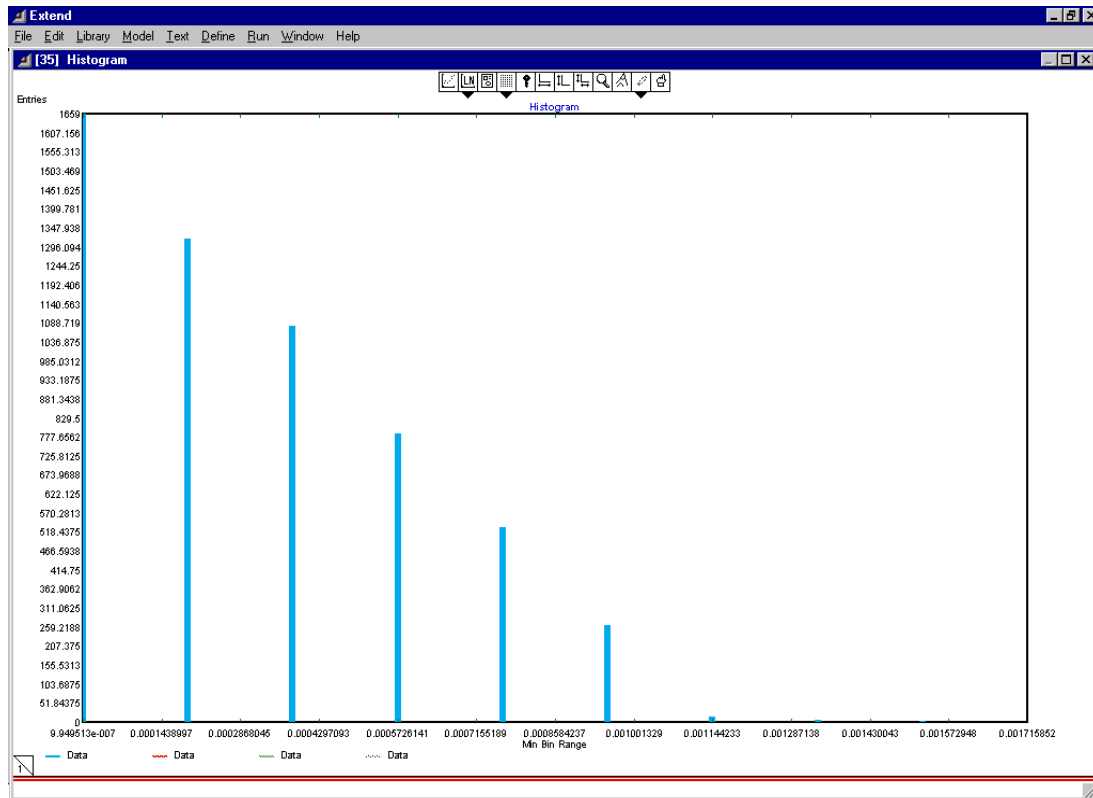


Figure 4: Histogram Representing Data for Product 21.

4 CONCLUSIONS

We implemented an integrated, real-time, in-line, dynamic output analysis procedure for controlling simulations of discrete manufacturing environments. We implemented this generic capability in Extend®, a commercially available simulation software from ImagineThat, Inc. We used this methodology to ensure that (1) time required for simulation and analysis was kept to a minimum, and (2) the analysts had an easy-to-use, relatively hands-off approach for ensuring that the results were statistically reliable. We are using this tool every day to effectively control simulations in electronics, steel, metal processing, and rubber industries.

ACKNOWLEDGMENTS

Many thanks to our colleagues at Maxager, Jim Collier, Vice President of Product Development, for encouraging and supporting this work, and Ted Emerson, Manufacturing Consultant, for his steel industry expertise.

REFERENCES

- Fishman, S.G., 1978. Grouping Observations in Digital Simulation. *Management Science*, Vol. 24, No. 3, pp. 510-521.
- Kelton, W.D., Sadowski, R.P, and Sadowski, D.A., 1998. *Simulation with Arena*, McGraw-Hill, New York, NY
- Lane, D.M., Testing Hypothesis using Standard Errors. *Hyperstat Online*. <<http://davidmlane.com/hyperstat/>>.
- Law, A.M., and Kelton, W.D., 1991. *Simulation Modeling and Analysis*, 2nd Edition, McGraw-Hill, New York, NY.
- Schmeiser, B., 1982. Batch Size Effects in the Analysis of Simulation Output. *Operations Research*, Vol. 30, No. 3, pp. 556-568.

AUTHOR BIOGRAPHIES

PETER MULLARKEY is a product architect at Maxager Technology, Inc. Prior to joining Maxager, he worked with Schlumberger from 1985-1999 doing research and development in knowledge-based decision support systems. Additionally, he worked for two years at Tandem Computers developing a state model-based diagnosis and maintenance system product. He received his undergraduate degree in Civil Engineering at Worcester Polytechnic Institute in Worcester, Massachusetts. He holds a Ph.D. in Civil Engineering from Carnegie Mellon University. His papers have appeared in AAAI and IEEE. His email address is <pmullarkey@maxager.com>.

SRINAGESH GAVIRNENI is a Software Architect at Maxager Technology, Inc. Previously he was a research Scientist in the Austin Product Center of Schlumberger. His undergraduate degree from IIT-Madras is in Mechanical Engineering and he holds Masters and Doctoral degrees in Manufacturing and Operations Management from Carnegie Mellon University. His research interests are in Supply Chain Management, Inventory Control, Scheduling, Mathematical Programming, and Simulation. His papers have appeared in *Management Science*, *Manufacturing & Service operations Management*, *European Journal of Operational Research*, *IIE Transactions*, and *IEEE Journal on Reliability*. His email address is <sgavirneni@maxager.com>.

DOUGLAS J. MORRICE is an Associate Professor in the MSIS Department at The University of Texas at Austin. He has a BA, Honours in Operations Research from Carleton University in Ottawa, Canada. He holds an M.S. and a Ph.D. in Operations Research and Industrial Engineering from Cornell University. His research interests include operations simulation design, modeling, and analysis. Dr. Morrice is a member of INFORMS. He served as the Secretary for the INFORMS College on Simulation (1994-1996) and was Co-Editor of the Proceedings of the 1996 Winter Simulation Conference. His email address is <morrice@mail.utexas.edu>.