SO YOU HAVE YOUR MODEL: WHAT TO DO NEXT A TUTORIAL ON SIMULATION OUTPUT ANALYSIS

Martha A. Centeno
M. Florencia Reyes

Industrial and Systems Engineering Department
Florida International University
Miami, Florida 33199, U.S.A.

ABSTRACT

Simulation models are built with the intent of studying the behavior of the real system represented by the model. However, a simulation model generates random outputs; thus, the data generated by it can only be used to *estimate* the true measure of performance. In this tutorial, we introduce several concepts and techniques to analyze such output. Additional examples will be given during the presentation of the tutorial.

1 INTRODUCTION

Simulation modeling enables the study of the stochastic behavior of systems, the testing of hypotheses that account for the observed behavior, and the use of these theories to predict future behavior. The simulation modeling methodology has several stages that begin with defining the objective of the study, model abstraction, model verification, and model validation. At the end of model validation, you finally have a working model. However, your job as a systems analyst is far from complete.

The actual analysis of the data provided by the simulation modeling will always depend on the initial objectives of the study, and on the type of system being modeled (Centeno, 1996; Sadowski, 1993). There are two main types of systems: *terminating* and *non-terminating*. Terminating systems have a natural *starting point* (operations begin) and a natural *ending point* (operations end), whereas non-terminating systems have a natural beginning, but they do not have a natural ending (Law, 1990). Typical terminating systems are most fast food restaurants, dental clinics, department stores, public-oriented government offices, and the stock market. There are other systems for which it is not obvious that they are terminating, but they are. For instance, a company producing the external tank for the space shuttle may be interested in studying several configurations to produce 12 tanks in 6 months. In this case, the system is terminating

because the natural ending of it is the production of the last external tank requested. Typical non-terminating systems are emergency rooms, some operations at hospitals, airports at large cities, and petrochemical plants. A manufacturing company that only works one shift may still be considered a non terminating system if the ending conditions for the shift are the initial conditions for the next shift (Law, 1990).

Depending on the type of objectives, the analysis may call for the comparison of various alternatives, or for the thorough analysis of the behavior of the systems under a specific configuration, or a quick analysis of a factor that may affect the performance of the system.

Every simulation model has various components, including *dynamic entities*, *resources*, and the *state of the system*. For the beneficiary of the simulation analysis, the important component is the state of the system because it is the collection of variables needed to describe the system's performance. An introduction to the analysis of these variables is the focus of this paper.

Section 2 discusses the importance of experimental design. A detailed example will be given when the tutorial is presented. Section 3 explains how to analyze terminating systems, whereas section 4 explains how to analyze non-terminating systems. Section 5 discusses how to compare several alternatives. Section 6 briefly presents what to do when dealing with a single replication. Finally, section 7 summarizes the tutorial.

2 EXPERIMENTAL DESIGN

When simulating a system, you have to think ahead about the scenarios of the system that you would like to evaluate using the simulation model (Kelton, 1994). In some cases, the scenarios are a natural consequence of the objectives and expectations set by management. In other cases, you need to find the right combination of the input parameters. If the number of inputs is small, exhaustive enumeration may be an easy and inexpensive way of finding the right

combination. But if the number of input factors is large, you should use some form of design of experiments to reduce the size of the search space (combinations to explore).

In the context of simulation, the *factors* are the various inputs to the model, the *levels* are the various options for each input parameter, and the *responses* are the outputs of the simulation model. Once these elements are clearly identified, you can analyze the experimental design in terms of measuring effects of the factors and the interactions among them. Donohue (1994) summarizes the various steps of using a statistical approach to the design and analysis of experiments, which also address tactical and strategic issues. A modified version of these steps is as follows:

1. Choose the factors (controllable input variables) and the response variables (uncontrollable output variables).
2. Define the region of operability (factor levels) and the regions of interest.
3. Select proper statistical analysis (ANOVA, regression, etc.), criteria to choose a "best" design, and appropriate experimental class design (e.g. factorial, Latin square, etc.)
4. Perform experiments, collect data, and analyze and summarize it. Incorporate here decisions regarding the type of system (terminating, non-terminating), simulation length, initial conditions, number of replications, and random seeds.
5. Draw inferences and conclusions

Your choices for the experiment design include complete factorials, fractional factorials, artificial factors, frequency domains, and correlated factors among others. Regardless of the chosen design, once you have collected the outputs, you must turn your attention to the response metamodels. The type of objectives set forth will dictate the selection of the appropriate metamodels (Donohue, 1994). At this stage, one may be interested in sensitivity analysis, in prediction, or in optimizing the response variables. Thus, we are interested in devising a model that characterizes the behavior of the system as a function of its outputs. Regression is commonly used to fit the behavior of the average of the performance measure and the variance of it. To add stability to the analysis, logarithmic transformations may be used (Sanchez, 1994).

Kelton (1994) provides an extensive list of references that give more details on designing the experiments. Sanchez (1994) and Donohue (1994) provide detailed examinations of the subject.

3 TERMINATING SYSTEMS

Terminating systems are systems that have a clear point in time when they start operations and a clear point in time when they end operations. For this type of systems, it is necessary to decide two things: the sample size and the simulation length. The simulation length is typically established by the context of the problem. For a car rental operation, it may be an entire day of operations, or it may just be the morning rush. The sample size is established based on the accuracy, reliability, and variation desired for the study, using the equation (1)

$$n = \frac{z_{\alpha/2}^2 \sigma^2}{d^2} \quad (1)$$

where d is the accuracy expressed in the same units as those of the measure of performance (e.g. within 2 unit), z is the critical value from the standard normal table at a given reliability level, $1 - \alpha$, (e.g. 95% reliability yields $\alpha = 0.05$), and σ is the standard deviation desired.

The resulting value of n is the minimum number of *replications* (not runs) needed to obtain statistically valid results. It is very common for the novice to confuse a replication with a simulation run. A *run* is what happens from the moment the user clicks on the run option of the main menu to the moment in which the software finishes outputting data and comes back to the main menu. A *replication*, on the other hand, is what happens from the simulated start time to the closing simulation time. In other words, *replication* is the repetition of a simulation with fixed inputs but different outputs due to different random numbers replications (Centeno, 1996; Clark, 1988). For a terminating system, a simulation run has n replications.

It is important to point out that n refers to the number of replications and not to the number of observations per replication. Why is this? Very simple! Assume that the measure of performance of interest is the time in the system (x_i), and that that you ran the simulation model so that 75 customers were processed. At the end of the replication run, there would be 75 values of x_i , one for each customer. These values would have been used to establish the replication's average time in the system ($\bar{x}_{within}$). $\bar{x}_{within}$ is an unbiased estimator of the true measure of performance's average (μ), but the variance of the observations within the replication (s_{within}^2) is biased because the x_i are not necessarily independent and identically distributed random variables; thus, it may happen that $\text{var}(\bar{x}_{within}) \neq \text{var}(x_i)/75$. To avoid this, we use the method of independent replications (Goldman, 1992). This method yields **one observation** per replication. So, at the end of the first replication where 75

customers were processed, you have one statistically independent observation of the time in the system ($y_i = \bar{x}_{within}$). Usually, terminating systems are analyzed using the method of independent replications. However, there are instances in which this method may be inefficient to use (see section 6).

Let us look at an example of the method of independent replications. Supposed that there is a driver licensing office that has seven officers to examine your driving skills. Customers arrive according to a normal distribution with mean 5 minutes and a standard deviation of 1 minute. It takes approximately 25 minutes to take the exam, normally distributed with standard deviation of 3.2 minutes. Further assume that you have generated a model of it, verify it and validate it. These are the things that you need to do now:

Establish the measures of performance for the analysis. From your objectives, you already know which measures of performance are important. Supposed that you are only interested on the average time in the system.

Decide the type of accuracy and reliability that you seek. Supposed that you want a 95% reliability ($\alpha = 0.05$), and an accuracy that yields an ideal half width confidence interval (h^*) equal to 3.

Run the model for a small number of replications. Depending on the size of the model, and the time it takes to execute it, small may mean 5, 10 or 15 replications. For this example, small means 10 replications. After running your model, you will obtain outputs similar to those given in Table 1.

Table 1: Output from a Terminating System

Replica- tion	Average Tsys	Replica- tion	Average Tsys
1	257.43	6	269.16
2	264.96	7	250.09
3	254.22	8	248.54
4	265.84	9	254.31
5	252.75	10	257.10

Compute a $(1-\alpha)\%$ confidence interval. This is done for the average of the measure of performance of interest, using equation (2).

$$\bar{x} - t_{n-1, 1-\alpha/2} \frac{s}{\sqrt{n}} \leq \mu \leq \bar{x} + t_{n-1, 1-\alpha/2} \frac{s}{\sqrt{n}} \quad (2)$$

where $t_{n-1, 1-\alpha/2}$ is the $(1-\alpha/2)$ percentile of the t-student distribution with $n-1$ degrees of freedom. The example data yields

$$\bar{x} = 253.44, s = 3.21, t_{10-1, 1-0.05/2} = 2.26$$

$$252.44 \leq \mu \leq 262.44$$

which means that the half width (h) is equal to 5. Recall that you are trying to answer the question of *how many replications to make, so that you obtain the appropriate resolution.* In other words, you are trying to find out if $h < h^*$. If after doing the preliminary run, the answer is **yes**, you are done. On the other hand, if the answer is **no**, you need to compute n^* using equation (3). This equation is the same as (1), but expressed in terms of the half widths of the confidence interval.

$$n^* = \text{Round} \left[n \times \left(\frac{h}{h^*} \right)^2 \right] = 28 \quad (3)$$

Modify simulation model to reflect the ideal n^* .

Change the number of replications, so that you can execute a full production run, and run the model.

Change the initial seed of the random number streams.

There are a couple of things that you need to decide before executing the production run.

You already ran 10 replications, but you need 28 replications. The questions that need to be answered are *Do you run the model again for 28 replications?*, or *Do you run the model for only 18 more replications?* If you run the model for 28 replications, the first 10 replications will be exactly the same 10 replications plus 18 new ones. This implies that you wasted some time. The first 10 replications are exactly the same because the random number streams begin exactly at the same point every time you run the model. So, to avoid this situation, what you need to do is to change the initial seed of the various random number streams that your model uses.

Compute the final $(1-\alpha)\%$ confidence interval as before. This yields the final confidence interval of the various measures of performance.

The method of independent replications can be used to build confidence intervals for statistics other than the mean value of the measure of performance. Other statistics of interests are the various percentiles. More details and examples can be found in Banks, Carson, and Nelson (1996).

4 NON TERMINATING SYTEMS

For non-terminating systems the fundamental question to answer is *for how long should the simulation be run?* To answer this question, you need two address two critical issues: 1) achieving steady state conditions, and 2) obtaining statistically independent observations. Once these questions are answered, it would be possible to obtain the confidence intervals as in the case of terminating systems.

As an example, consider the processing of the space shuttle which requires a large number of assembly and repair operations. Technicians working on the shuttle receive the necessary assembly pieces and tools in kits. These kits are customized for the various operations, but

fall into 10 different categories. Suppose that request for assembly kits come to the kitting shop according to a triangular distribution with mode 20 minutes, and range (12, 25). The shop has 4 technicians who assemble the elements of the kit on cardboard. Once the cardboard is ready, it is placed on a conveyor that transports it to a wrapping machine. It takes a technician approximately 15 minutes to assemble a request (according to an exponential distribution), and it takes the wrapping machine 30 seconds to press the wrap onto the kit. Traversing the conveyor belt takes approximately between 2 and 3 minutes. Assume that you have written, verified, and validated your model. These are the things that you need to do now.

Establish the measures of performance for the analysis. Again, this comes from the objectives of the study. Supposed that you are only interested in the average time it takes to fulfill an order (time in the system).

Decide the type of accuracy and reliability that you seek. Supposed that you want a 95% reliability ($\alpha = 0.05$), and an accuracy that yields an ideal half width confidence interval (h^*) equal to 3.

Run the model for a short simulation length. Depending on the size of the model, and the time it takes to execute it, a short run may mean 500 time units or 5,000 or 10,000. For this example, short means 10,000 time units. Make sure that you save to a file the individual observations of the time in the system. These values will be used later on.

Establish the warm up period. The warm-up period is that period in the life of non-terminating systems in which the system was merely filling up. Typically, data generated during this period instills bias in the analysis. This phenomenon is known as the bias of the initial conditions. Some of the proposed ways of getting rid of the bias include the truncation of the output data and making a very long simulation run. In the first case, all data generated during the warm-up time is eliminated from the final analysis. In the second case, nothing gets eliminated because the theory goes that if the simulation length is sufficiently large, then the initial conditions have minimal, if any, effect on the steady state behavior of the system.

Deciding on the length of the warm up period has been the subject of intense study. Several researchers have proposed a diversity of methods. See Goldsman (1992) for a list of these efforts. To illustrate one of the methods, we will use the moving average approach in conjunction with visualization of the data.

From the example, we can generate a graph similar to that in Figure 1. By looking at the graph, the fill time seems erratic from entity to entity (individual times); however, they seem to gather mostly between 12 and 24 time units. Thus, what we need to do is to see if by collecting them in small groups, we can identify the trend that is follows, if any. Moving averages do precisely that.

Figure 1: Individual Observations and Moving Average

To establish the warm up time, we look at the moving averages. From Figure 2, it seems that the initial conditions do not have much impact on the behavior of the fill time after 3,000 units of time. Thereafter, it is as random as it will ever be! So, the warm up (WT) period time will be 3,000 time units.

Figure 2: Moving Average

Create batches. Since the individual observations within a replication are not independent, we cannot compute a CI for the average of the measure of performance using the individual data points. To overcome this, we use the method of batch means. This approach divides the individual observations of a single long simulation run into a number of contiguous batches. Once these batches are formed, we can invoke the benefits of the central limit theorem because each batch will hopefully yield one iid observation.

To establish these batches, we need to know the batch size and how many batches are needed. To determine the batch size, we need to find the point in the set of observations where any two observations (k observations apart from each other) have a minimum correlation ($\rho \rightarrow 0$). The value of k is known as the *lag* between two observations. In the case of our example (Figure 3), $\rho \approx 0$ when the lag = 90; thus, a reasonable batch size is 10 times

the value of that lag (900 observations per batch). The number 10 is an empirical value.

Figure 3: Correlogram

To establish the number of batches, simply guide yourself by the requirements of building confidence intervals: the more the merrier. You should never go below 10 batches. Values of above 30 or more batches are preferred. Let us say that for our example we settled for 20 batches. So having the need of generating data for 20 batches means the following:

$$900 \frac{\text{observations}}{\text{batch}} * 20 \text{ batches} = 18,000 \text{ observations}$$

Based on the average inter-arrival time, we get 12 observations / hour. Thus, to compute the generation time (GT), we use

$$GT = \frac{18,000 \text{ observations}}{12 \text{ observations / hour}} * \frac{60 \text{ min}}{1 \text{ hour}} = 90,000 \text{ min}$$

Finally, the simulation length (SL) is as given by:

$$SL = WT + GT + \text{safety time} \quad (4)$$

The safety time is an added protection in search of eliminating the effect of the initial conditions. A rule of thumb is to use 1% of GT.

Construct confidence intervals. Once the batches are generated, each batch takes the place of a replication. Therefore, we can use the methods utilized for terminating systems from this point forward.

When using the method of batch means, it is important that you pay attention to the covariance between batches. Having small correlation is a necessary but not sufficient condition for independence; the covariance can further strengthen or weaken the independence assertion: the smaller the covariance, the better. To get it smaller, eliminate as much as you can of the “initial” observations.

As long as you maintain at least 10 batches, you are relatively OK.

5 COMPARING ALTERNATE SYSTEMS

The greatest benefit of using simulation modeling is the ability to compare different system configurations before deciding which systems to implement. A critical issue in this case is how the various alternatives are run. Because simulation uses random number streams, different streams will lead to different results. One must be able to assess if the observed difference is due to the proposed system configuration or due to the effect of the random number streams. Several approaches have been proposed to deal with this issue. The use of common random numbers is strongly suggested when comparing several alternate systems. The main idea is to use the same set of random numbers for each one of the alternatives, so that there is a large positive correlation among the values of their performance measures; thus reducing the variance of the pair-wise difference between the alternatives (Clark, 1988).

Any two systems can be statistically compared through a variety of procedures such as single paired-t confidence intervals, and two sample t-confidence intervals. The most appealing method is the single paired-t because:

- ✓ It does not require that the variance of the measure of performance under consideration (for both systems) be equal (or assumed equal).
- ✓ It does not require that the observations from system 1 be independent of the observations from system 2 (or vice versa), as long as the observations from each system are independent within the system.
- ✓ It reduces the problem of comparing two systems to the estimation of a single parameter (the difference); thus, reducing the statistical bias.
- ✓ It is the same as the hypothesis testing case where the mean (of a measure of performance) of two system is to be tested for any differences, using the individual observations of each sample instead of a summary (single value) from each sample.

These are the steps that you need for using the single paired-t method.

Collect data from the two alternate models of the system making sure that a big enough sample is gathered, and that the same number of observations are taken from both systems (if not, you may need to conduct a full ANOVA test with incomplete data points).

Compute the differences.

$$d_i = x_{1i} - x_{2i}$$

where d_i = individual differences.

x_{1i} = i th observation from system 1.

x_{2i} = i th observation from system 2.

d_i is a m. v. u. e. of δ .

δ is the parameter that represents the true difference between the two systems.

Compute the average difference and its standard deviation.

$$\bar{d} = \frac{\sum d_i}{n} \quad s_{\bar{d}} = \frac{s_{d_i}}{\sqrt{n}} \quad \hat{\delta} = \bar{d}$$

Compute the confidence interval around $\bar{d}$.

Interpret results. If the interval contains the value 0, nothing can be said about the systems; they may be the same or they may be different; therefore, run more replications. If the interval does not contain the value 0, then it can be stated that both systems are different with $(1-\alpha)$ confidence. The meaning of $\hat{\delta} > 0$ or $\hat{\delta} < 0$ must be interpreted based on the specific measure of performance and problem statement. As an example consider the output in Table 2. Assume that the measure of performance is time in the system. System 1 has 1 server, whereas system 2 has 2 servers.

Table 2: Sample Output

replication	time in the system - system 1	time in the system - system 2	difference
1	3.4139	1.0787	2.3352
2	3.8983	1.0988	2.7995
3	2.2599	1.0943	1.1656
4	1.9491	1.0578	0.8913
5	3.4604	1.0963	2.3641
6	2.7795	1.1004	1.6791
7	3.2541	1.1040	2.1501
8	3.5396	1.0966	2.4430
9	4.0727	1.1460	2.9267
10	2.4621	1.0771	1.3850

$$\bar{d} = 2.0140 \quad s(d_i) = 0.6608 \quad t_{10-1, 1-0.05/2} = 2.26$$

$$1.541 \leq \delta \leq 2.487$$

What can you say about it? In this particular example, since we are comparing time in the system, and it is assume that a shorter time in the system is "better", then

from the confidence interval, we can say that the configuration for system 2 yields a better performance.

6 SINGLE REPLICATION SIMULATION

In many cases, simulation may be used for a "quick and dirty" examination of the effect of changing the system's environment. In these instances, a detailed simulation analysis may not be feasible; in fact, a single short replication may be sufficient to achieve the desired goal. Single replication simulations are very useful when simulation is used for real time decision-making. For example, given that an event (rare but significant) has occurred, one may wish to simulate the system to determine how should the operations of the system be changed to maintain the level of performance.

In *single-replication* simulation, random variation must be carefully controlled. One way to do this is to treat infrequent, yet important, events as the initial conditions of a simulation; i.e. the rare events are pulled out of the simulation model run (Goldsman, Swain and Withers, 1990). In addition, one may control the sources of randomness by altering the way sampling occurs when the simulation is executed. Goldsman, Swain, and Withers (1994) have proposed the truncation of the sampling distributions, conditional sampling, and selective alteration of the distribution parameters. In the end, you will have a single observation of the statistic desired.

Point estimation provides minimal information. Most of the time, we are interested in obtaining a confidence interval for the measure of interest. However, having only a single replication eliminates the possibility of using the classical z- or t- methods for obtaining confidence intervals. It has been proposed that some Bayesian methods may be used to obtain the confidence intervals. But, these methods require that we establish parameters about the population, which is indeed unknown. Thus, these methods should be used with caution. Nonetheless, these methods will enable the analyst to better understand the performance of the systems even if only one replication is done.

7 SUMMARY

You have been introduced to several of the many aspects of analyzing simulation outputs. You are strongly encouraged to further review the literature on this subject, in particular where it pertains to design of experiments and steady state analysis. For other methods and more details of steady state analysis, see the works by Charnes (1993), Heidelberger and Welch (1983), and Chance and Schruben (1992). A valuable approach for simulation outputs analysis is graphical analysis. For details on this approach see the works by Grier (1992) and Law and Kelton (1991).

An important concept that you should keep in mind is that the raw output from simulation models is not independent and identically distributed. Whether the system is terminating or non-terminating, you should use a technique that will transform the raw output into iid observations.

Also, you should keep in mind that simulation software is continuously evolving to provide support with output analysis. Several commercially available packages already incorporate some form of support for the analysis of simulation outputs.

REFERENCES

- Banks, J., J. S. Carson, II, and B. L. Nelson. 1996. *Discrete-Event System Simulation*. Upper Saddle River, New Jersey: Prentice Hall.
- Centeno, M. A. 1996. An Introduction to Simulation Modeling. In *Proceedings of the 1996 Winter Simulation Conference*, ed. J. M. Charnes, D. J. Morrice, D. T. Brunner, and J. J. Swain, 15-22.
- Chance, F. and L. W. Schruben. 1992. Establishing a Truncation Point in Simulation Output. Technical Report, School of Operations Research and Industrial Engineering, Cornell University, Ithaca, New York.
- Charnes, J. M. 1993. Statistical Analysis of Output Processes. In *Proceedings of the 1993 Winter Simulation Conference*, ed. G. W. Evans, M. Mollaghasemi, E. C. Russell, W. E. Biles, 43-48.
- Clark, G. M. 1988. Tutorial: Analysis of Simulation Outputs to Compare Alternatives. In *Proceedings of the 1988 Winter Simulation Conference*, ed. M. Abrams, P. Haigh, J. Comfort, 19-23.
- Goldsman, D. 1992. Simulation Output Analysis. In *Proceedings of the 1992 Winter Simulation Conference*, ed. J. J. Swain, D. Goldsman, R. C. Crain, and J. R. Wilson, 97-103.
- Goldsman, D., J. J. Swain, and D. Withers. 1990. Single Replication Simulation. In *Proceedings of the 1990 Winter Simulation Conference*, ed. O. Balci, R. P. Sadowski, R. E. Nance, 387-391.
- Grier, D. A. 1992. Graphical Techniques for Output Analysis. In *Proceedings of the 1992 Winter Simulation Conference*, ed. J. J. Swain, D. Goldsman, R. C. Crain, and J. R. Wilson, 314-319.
- Heidelberger, P. and P. D. Welch. 1983. Simulation Run Length Control in the Presence of an initial transient. *Operations Research*, 32, 185-209.
- Kelton, W. D. 1994. Analysis of Output Data. In *Proceedings of the 1994 Winter Simulation Conference*, ed. J. D. Tew, S. Manivannan, D. A. Sadowski, and A. F. Seila, 62-68.
- Law, A. M. 1990. Design and Analysis of Simulation Experiments for Manufacturing Applications. In *Proceedings of the 1990 Winter Simulation Conference*, ed. O. Balci, R. P. Sadowski, R. E. Nance, 33-37.
- Law, A. M. and W. D. Kelton. 1991. *Simulation Modelling and Analysis*, McGraw-Hill, New York.
- Sadowski, R. 1993. Selling Simulation and Simulation Results. In *Proceedings of the 1993 Winter Simulation Conference*, ed. G. W. Evans, M. Mollaghasemi, E. C. Russell, W. E. Biles, 65-68.
- Sanchez, S. M. 1994. A Robust Design Tutorial. . In *Proceedings of the 1994 Winter Simulation Conference*, ed. J. D. Tew, S. Manivannan, D. A. Sadowski, and A. F. Seila, 106-113.
- Seila, A. F. 1990. Output Analysis for Simulation. In *Proceedings of the 1990 Winter Simulation Conference*, ed. O. Balci, R. P. Sadowski, R. E. Nance, 49-53.
- Seila, A. F. 1991. Output Analysis for Simulation. In *Proceedings of the 1991 Winter Simulation Conference*, ed. B. L. Nelson, W. D. Kelton, G. M. Clark, 28-36.

AUTHOR BIOGRAPHIES

MARTHA A. CENTENO is an associate professor in the Department of Industrial and Systems Engineering at Florida International University. She received a B.S. degree in Chemical Engineering from ITESO University (Mexico), an M.S. degree in Industrial Engineering from Louisiana State University in 1985, and a Ph.D. degree in Industrial Engineering from Texas A&M University in 1990. Her research interests are in the areas of intelligent simulation, smart data analysis, and engineering education. Dr. Centeno is a member of Alpha Pi Mu, Tau Beta Pi, Phi Eta Sigma, IIE, INFORMS, SCS, ASQ, ASA, and ASEE.

M. FLORENCIA REYES received a B.S. degree in Industrial and Systems Engineering from Florida International University in 1997. She is completing her M.S. degree in Industrial Engineering from Florida International University. Her research interests are in the areas of goal driven simulation. She is a member of Alpha Pi Mu, Tau Beta Pi, Golden Key, Phi Eta Sigma, IIE, and INFORMS.