Proceedings of the 1999 Winter Simulation Conference
P. A. Farrington, H. B. Nembhard, D. T. Sturrock, and G. W. Evans, eds.

HEALTHCARE SIMULATION: A CASE STUDY AT A LOCAL CLINIC

Mark L. Weng
Ali A. Houshmand

Department of Industrial Engineering
University of Cincinnati
Cincinnati, OH 45221-0072, U.S.A.

ABSTRACT

Today, researchers and analysts are beginning to uncover the potential for using simulation in the health care field; with a multitude of interactions between patients, physicians, nurses, and technical and support staff, simulation can be an invaluable tool. Inefficiencies can be eliminated or resource allocation changed to determine an optimal setup. Primarily, simulation has been used in the health care field in comparison studies of alternative systems for resource or scheduling requirements (Lowery 1998). When analyzing such alternatives, the standard performance measures are typically reported: throughput, time in system, and queue times and lengths. This paper is a systems analysis of a clinic using the above mentioned performance measures along with another proposed performance measure, total cash flow.

1 INTRODUCTION

A local health care administration group that runs several hospitals in the Cincinnati area began work on a project to analyze one of their clinics to determine if it can become more profitable. Management has already initiated the data collection process and data analysis, and they have concluded that there are multitudes of inefficiencies in the system causing the reduced profitability. Realizing that simulation could be a useful tool for analyzing the clinic, the health care management group sponsored a project under the direction of Dr. Ali Houshmand at the University of Cincinnati to model and to analyze the clinic.

Meetings with management indicated that they wanted a base model that captured the essence of the system without sacrificing accuracy. Once the base model was completed, a series of scenarios would be provided for analysis, with the key performance measures being:

- To maximize patient throughput, and
- To reduce patient time in the system.

The solution to these two problems is one that is not easily determined. From an engineering standpoint, the solution is not too difficult to find. It is the human factors standpoint that makes determining a solution very tricky. Due to the complexity of the clinic interactions and the management versus doctor issues involved, it has been virtually impossible to optimize the system; management wants to maximize throughput and reduce the time patients spend in the clinic while doctors want to provide quality health care for patients and maintain a comfortable level of autonomy, something that has been taken from them since the introduction of managed care. Without compromising either party's interests, changes in the system need to be made to determine if it is at all possible to reach a middle ground that satisfies both parties involved.

Since the power of simulation lies in its abilities to model alternative systems for comparison studies and estimate a number of varying performance measures, simulation was a natural solution.

2 DESCRIPTION OF THE CLINIC

The clinic being modeled is considered an outpatient clinic at a local hospital. Several different patient classifications visit the clinic; there are urgent care and acute care patients, with the majority being acute care. Most of the urgent care patients are sent to the hospital emergency room. Furthermore, there are return patients, i.e., patients with appointments, and new patients, i.e., walk-ins. A flow chart of a typical patient's process would be similar to the following: A patient arrives in the clinic and goes to the waiting room for registration. The registrar calls the patient to the registration booth for registration. After registration is complete, the patient returns to the waiting room to wait for the medical assistant to call him to begin his consultation. The medical assistant will call the patient back to take his weight on a scale and his blood pressure in the doctor's exam room. Once the doctor is ready, he goes to the exam room and begins the doctor consultation. Usually more than halfway through the consultation, the doctor usually has to meet with the preceptor (doctor in charge) for advice and to verify his diagnosis. Following the meeting with the preceptor, the doctor returns to the patient, still in the exam room, for the final diagnosis. He may order a specific lab test, in which case the patient proceeds to the lab for the respective test, or the patient can leave. In either case, the consultation is over, and the patient must proceed to the receptionist to schedule a follow-up visit. This is the standard process that was modeled; however, there are variations on this process that will not be addressed.

3 DATA COLLECTION

The data collected by the benchmarking engineers over a period of two months were provided for use in this project. These data mainly consisted of patient process times and patient arrival times. An interview with the staff manager provided the standard staff scheduling times and requirements as well as the probabilities for process branches. Data collected on the clinic expenses based on 16,188 cases for fiscal year 1998 provided such information as total revenue and expense data for the variable and fixed labor, variable and fixed supplies, overhead, depreciation, and medical education. Furthermore, a breakdown on the average salaries of the staff included the following:

- Medical Assistant
- Preceptor
- Registrar
- Receptionist
- R1 (First-year residents)
- R2, R3 (Second-, Third-year residents)

Since management is interested only in analyzing the effects of the R1's and the remaining residents (R2's and higher), the salaries for the R2's and R3's were averaged between the two groups.

4 MODELING ASSUMPTIONS

In order to eliminate any possible misunderstandings or unrealized expectations, a list of the modeling assumptions were provided to the management team for review. These assumptions are stated as follows:

- *The patient arrival rate is modeled as a nonstationary Poisson process.*

  Since the arrival rate is not steady throughout the day, with more patients arriving at the clinic opening and just after noontime than in the other hours, a nonstationary Poisson arrival rate is appropriate. Law and Kelton (1991) and Nelson (1995) provide insightful descriptions and analyses of nonstationary Poisson processes, and Kelton, Sadowski, and Sadowski (1998) describe a modeling technique called thinning used in this model.
- *The duration of the simulation is 600 minutes, from 8AM to 6PM consisting of two five-hour sessions.*

  The clinic generally operates between 8AM and 5PM, but the additional time is to allow the clinic to clear the patients out of the system. The clinic does not close until the last patient leaves. The clinic operating hours includes two separate sessions, the first from 8AM to 1PM and the second from 1PM to 6PM.
- *The base model contains an "average" staff.*

  The base model will use what the clinic's staff director defined as the standard staffing requirements listed in Table 1 below.

Table 1: Resource Requirements used in Simulation Model

RESOURCE	**Number in Simulation Model**
Receptionist	2
Registration Clerk	2
Medical Assistant	4
R1	4
R2	4
Preceptor	2

- *New and return patients have different process times for various processes.*

  New and return patients have different service requirements and times. Therefore, those times are used where applicable.
- *The clinic is modeled as a pull system.*

  In a pull system, entity flow is created by demand generated by downstream processes. After entering the clinic and completing registration, patients must wait until an exam room and their doctor is available. Only at that time can they leave the waiting room to begin their consultation. It is erroneous to model the clinic as a push system, where the entity flow is created by introducing entities into the system. To model a valid pull system, patients cannot leave the waiting room unless a medical assistant is available.
- *Expense data in the model are based on hourly averages.*

  In order to ensure that the yearly salaries remain constant regardless of the simulated time, the cash flow portion of the model uses data converted into hourly averages instead of basing them by case.

5 MODEL CONSTRUCTION

5.1 Selection of Software

The simulation language, Arena, was used to model the clinic. Arena is a Windows-based platform that is popular and widely used due to its tremendous flexibility and ease of use. Models can be constructed without any programming knowledge due to its use of dialog boxes and graphical interface, but low-level programming can be done if so desired. Arena has been used to model everything from manufacturing work cells to complex interactions in an emergency room.

Furthermore, Arena's Input and Output Analyzers provide excellent tools to fit input probability distributions based on actual data and analyze output data using classical statistical measures. Animation of the clinic provides a visual representation of the clinic for those without the technical understanding of a simulation modeling language. For a better understanding of Arena, Kelton, Sadowski, and Sadowski's (1998) book provides an excellent overview of simulation using Arena for beginners and experts alike.

5.2 Input Distributions

The clinical processes are stochastic in nature, so discrete data cannot be used in its present form. In order to model a stochastic process with random inputs, probability distributions need to be specified. Therefore, the data needs to be fitted to an input distribution for Arena to utilize. Since there were very few overall observations, some of the resultant input distributions reflect the sparseness of data; there were some triangular and uniform distributions fitted which were verified through expert knowledge, i.e. questioning the clinic's staff manager. However, the majority of input distributions fitted for the processes were gamma and beta distributions, which are representative of time required to complete a task and absence of data. The input distributions used in the Arena model were all fitted according to Arena's Input Analyzer. Table 2 shows the processes and their respective input distributions. See Law and Kelton (1991), Banks, Carson, and Nelson (1996), and Nelson (1995) for more on fitting input distributions.

Table 2: Process Time Data and Distributions

PROCESS	RESIDENT	DISTRIBUTION
Registration	NA	2 + 21 * BETA(0.829, 1.09)
Weighing	NA	2.77 * BETA(1.39, 3.71)
Vitals	NA	LOGN(7.25, 9.47)
Initial Diagnosis	R1 – new	TRIA(29, 58, 87)
	R1 – return	13 + ERLA(10.4, 1)
	R2 – new	TRIA(14, 35.9, 87)
	R2 – return	5 + 31 * BETA(0.846, 0.627)
Write Patient Records	R1	GAMM(39.7, 0.589)
	R2	1 + 8 * BETA(0.668, 1.06)
Consult Preceptor	R1 – new	5 + LOGN(6.17, 22)
	R1 – return	1 + GAMM(12.6, 1.15)
	R2 – new	UNIF(3.57, 7)
	R2 – return	3 + ERLA(3.44, 1)
Complete Records	R1	0.06 + GAMM(1.02, 2.27)
	R2	9 * BETA(0.733, 1.19)
Final Diagnosis	R1 – new	2 + 17 * BETA(0.405, 0.541)
	R1 – return	LOGN(8.24, 21.8)
	R2 – new	1 + 18 * BETA(0.535, 0.761)
	R2 – return	TRIA(0, 0.112, 14)
Lab Test	NA	6 + 34 * BETA(0.8, 1.55)
Finish Paperwork	NA	GAMM(3.7, 1.76)

5.3 Special Modeling Techniques

Perhaps the most important feature of the clinic is the fact that it is a pull system. To model this aspect of the clinic, patients are duplicated following registration, with the original sent to a queue where it must wait for a signal before it can proceed further along in the model, and the other duplicate sent along a separate signal branch. In this signal branch, the model checks for an available medical assistant, since the medical assistant must initiate the patient consultation process. If the medical assistant is available, a signal is sent to the waiting patient entity in the main model that allows the original patient to continue through the system. Meanwhile, the duplicate is disposed from the entire model.

To incorporate cash flow into the system, a separate submodel was constructed that has the variable expense entities arriving into the submodel every 60 minutes. The accounting data were broken down from average revenue and expense per year to average revenue per hour and variable expense per hour. A variable, called VExpenses, sums the variable expenses per hour each hour for each process of the clinic utilizing variable expenses. At the end of the simulated time, a final value for the VExpenses is displayed. Furthermore, a time-series graph in the model displays how the VExpenses variable changes over time. Similar to the animation, this allows one to simply view the graph to understand how variable expense cash flow is changing over time. In the same way, a variable called Revenues was created to aggregate the revenues for the patients entering the system. Likewise, a graph indicating the change in the total revenue is provided. Thus, in order to determine the final profit, the Vexpenses value is subtracted from the Revenues value and the fixed expenses calculated from the data. The fixed expenses value is also subtracted from the Revenues value to yield a final value indicating the profitability of the clinic.

As stated in the assumptions, the base simulation time is one business day, ten hours. This is to reproduce what one could expect to observe for any given day. For statistical validity, i.e., reducing the variance, a large number of replications needed to be made. Examination of the confidence interval half-widths for the mean simulated values indicates less than five-percent error. An in-depth discussion on model validity follows in the next section. Since the clinic system has a definite starting and stopping time (terminating system), the model does not require any warm-up time.

6 MODEL VALIDATION

To validate the system aspects of the model, several measures have been taken. First, the model includes animation, and the patients proceed along their respective paths. Nothing occurs out of the ordinary; entities move to the locations where they are expected to go. Second, after running 60 replications, a quick statistical analysis comparing the actual average initial and final diagnosis times for both the R1's and the R2's to the results that Arena generated was made. The comparisons along with a percent error are shown in Table 3. As can be seen in the table, the percent error for the R1's is very low, less than one percent. However, the percent error for the R2's is much higher, 6.5% for the initial diagnosis, 22.1% for the final diagnosis, and 8.7% for the total patient interaction time. This numerical disparity is likely due to the nature of using percent errors. Since the numbers for the R2's were relatively low, any deviation from those numbers would produce a percent error greater than the same deviation from a larger number. Furthermore, due to the paucity of data, the input distributions cannot be guaranteed to be correct. Collection of additional data may prove that the fitted input distributions are not accurate.

A more accurate measure would be to show where the actual data lies in the confidence intervals provided by Arena. If the actual mean time lies in the 95% confidence interval that Arena's Output Analyzer generates, then one can conclude that in the long run, the mean actual times will lie within the confidence interval 95% of the time. The confidence intervals are displayed in Table 4.

Table 3: Comparison of Simulation Results to Actual Data

	ACTUAL		**SIMULATION (60 Replications)**		**PERCENT ERROR (%)**	
	R1	*R2*	*R1*	*R2*	*R1*	*R2*
Initial Diagnosis	34.51	26.67	34.20	28.40	0.89	6.50
Final Diagnosis	7.48	4.32	7.51	5.28	0.38	22.11
Total Patient Interaction Time	41.99	31.00	41.71	33.68	0.67	8.68

Table 4: Actual Diagnosis Times with Arena Confidence Intervals

	ACTUAL R1	**95% CI**	**ACTUAL R2**	**95% CI**
Initial Diagnosis	34.507	(31.9, 36.5)	26.667	(26.6, 30.3)
Final Diagnosis	7.482	(6.59, 8.42)	4.324	(4.91, 5.65)

With the exception of the R2 final diagnosis, the actual times fall within the 95% confidence intervals given by Arena. This indicates that the consultation time is well-modeled. However, to determine the accuracy of the entire model, the essential aspect of the system that needs to be analyzed is the total aggregate flowtime through the clinic; it is necessary to compare Arena's results with the actual data in order to validate the entire model. For 60 replications, Arena obtained an average patient flowtime of 95.2 minutes with a 95% confidence interval of (92.8, 97.5). The actual average patient flowtime is 97.41 minutes. Therefore, we can conclude that the simulation accurately models the entire patient process.

7 ANALYSIS

7.1 Scenario Analysis and Comparison

As mentioned at the beginning of this paper, the health care management group is interested in maximizing the patient throughput and reducing total time in the system. With all variables remaining constant and varying only the number and types of residents, they would like to know which of the following three scenarios maximizes throughput and minimizes total time in system:

- Scenario 1: 8 R2's only,
- Scenario 2: 6 R2's only, and
- Scenario 3: 8 R1's only.

A statistical analysis of the throughput and total time in system will help determine what the optimal alternatives are. Arena's Output Analyzer can run a paired-t test on data gathered across the 60 replications for comparing any two scenarios. After running 60 replications for each scenario, the Arena confidence intervals for the patient throughput is shown in Table 5. It is evident that the different scenarios provide different mean flowtimes, but in order to determine if there is a statistical difference, a paired-t test needs to be performed. Under the assumption that the variances are equal, the paired-t test is appropriate. The Output Analyzer also includes a two-sample t-test for comparing two means, but since there was no assumption of statistical independence, the paired-t test is the correct test to use. For all three scenarios, there is a statistical difference between the base model and each scenario. The hypothesis tests rejected the null hypothesis that there is no difference between the two means. The two scenarios featuring only the R2's resulted in significantly lower flowtimes than the base model, and the third scenario with all R1's resulted in a significantly higher flowtime than the base model. This leads to a very important finding: there is a significant difference between the R1 consultation time and the R2 consultation time, which ultimately leads to a significant difference in patient flowtimes. Therefore, to minimize patient flowtime, R2's should only be used. However, using only R2's is impossible; R1's still need to gain experience. The next best alternative is to separate the R1's and the R2's; optimize scheduling by using all R2's and all R1's when possible. The next step is to compare the scenarios of 8 R2's versus 6 R2's. Is there really difference between the two systems when, all else constant, the number of R2's varies by two? A paired-t comparison between the mean patient flowtime with 8 R2's and with 6 R2's indicates that there is no significant difference between the two scenarios. Before jumping to any conclusions, it must be noted that patient throughput needs to be considered in this instance. When comparing the mean patient throughput for Scenario 1 (20.8 patients) against the mean patient throughput for Scenario 2 (20.3 patients), Scenario 1 yielded 0.5 more patients than Scenario 2. The paired-t test indicated that there is no statistical difference between the two scenarios. Therefore, we can conclude that Scenario 2 (6 R2's) is optimal.

Table 5: Comparison of Mean Simulated Flowtime

	BASE	**8 R2**	**6 R2**	**8 R1**
Time in Minutes	101	88.2	89	107
Half-Width	3.35	3.02	3.07	3.5
Half-Width Error	3.32%	3.40%	3.45%	3.27%
95% Confidence Interval	(98.1, 105)	(85.8, 91.8)	(85.9, 92.1)	(103, 110)

Table 6: Baseline Model Simulation Summary

AVERAGES	VARIABLE EXPENSES	REVENUE	PROFIT	PATIENTS SEEN	TOTAL PATIENTS IN SYSTEM	PERCENTAGE
Baseline	-$1064.83	$1377.96	$313.13	17.5833	18.0667	97.33%
Scenario 1	-$798.60	$1216.01	$417.41	15.5167	17.7667	87.34%
Scenario 2	-$800.23	$1247.35	$447.12	15.9167	18.4167	86.43%
Scenario 3	-$816.53	$1253.89	$437.35	16.0000	18.5000	86.49%

7.2 Scenario Analysis and Comparison with Expenses

7.2.1 Baseline Model

Next, the scenarios are evaluated using cash flow as the performance measure. Using the baseline simulation model for each of the three scenarios along with the default setup, Table 6 below summarizes the results. The values are all results for one business day.

The fixed costs for each day total to the following:

Fixed Labor:	$784.22
Fixed Supplies:	$53.26
Fixed Miscellaneous:	$2000.76
Depreciation:	$20.97
Department Overhead:	$780.03
Hospital Overhead:	$2555.67
Medical Education:	$365.60
Preceptors (2):	$769.24
Total:	$6560.51

However, fixed labor costs breakdown according to the staffing sizes.

Baseline:
- 4 R1's: $128.08*4 = $512.32
- 4 R2's: $133.85*4 = $535.40

Total: $1047.72

Scenario 1:
- 8 R2's: $133.85*8 = $1070.80

Scenario 2:
- 6 R2's: $133.85*6 = $803.10

Scenario 3:
- 8 R1's: $128.08*8 = $1024.64

Therefore, the final dollar figures for the baseline models are:

Baseline:	-$7295.10
Scenario 1:	-$7213.90
Scenario 2:	-$6916.49
Scenario 3:	-$7147.79

The smallest loss would be for Scenario 2 with only 6 R2's. From the table above, 86.425% of the patients that show up in the system are cleared out before 5 PM in the simulation model. This is the lowest of the three scenarios, but considering that there are two fewer residents than in either Scenario 1 or 3, this is an acceptable level. These results agree with the results obtained from investigation of the patient flowtime performance measure.

7.2.2 New Standards Model

A similar analysis can be done for the scenarios based on the revised standards proposed by the health care management benchmarking group. The results are summarized in Table 7.

Table 7: New Standard Model Simulation Summary

AVERAGES	VARIABLE EXPENSES	REVENUE	PROFIT	PATIENTS SEEN	TOTAL PATIENTS IN SYSTEM	PERCENTAGE
New Standard	-$840.81	$1280.00	$439.20	16.3333	17.8000	91.76%
Scenario 1	-$829.27	$1324.41	$495.14	16.9000	18.5167	91.27%
Scenario 2	-$829.27	$1324.41	$495.14	16.9000	18.5167	91.27%
Scenario 3	-$836.00	$1270.86	$434.86	16.2167	18.3667	88.29%

The fixed costs for each day total to the following:

Fixed Labor:	$784.22
Fixed Supplies:	$53.26
Fixed Miscellaneous:	$2000.76
Depreciation:	$20.97
Department Overhead:	$780.03
Hospital Overhead:	$2555.67
Medical Education:	$365.60
Preceptors (2):	$769.24
Nurses (2):	$384.62
Total:	$6945.13

However, fixed labor costs breakdown according to the staffing sizes.

New Standard:
- 4 R1's: $128.08*4 = $512.32
- 4 R2's: $133.85*4 = $535.40

Total: $1047.72

Scenario 1:
- 8 R2's: $133.85*8 = $1070.80

Scenario 2:

- 6 R2's: $133.85*6 = $803.10

Scenario 3:

- 8 R1's: $128.08*8 = $1024.64

Therefore, the final dollar figures for the new standard models are:

New Standard:	-$7553.66
Scenario 1:	-$7520.79
Scenario 2:	-$7253.09
Scenario 3:	-$7534.91

Once again, the smallest loss would be for Scenario 2, at -$7253.09 per day. In addition, the percentage of patients seen is the highest after the new standard model with four medical assistants instead of the two in the second scenario. These results make sense, since any way to see a high number of patients while reducing any costs should result in the lowest lost profit. Furthermore, the proximity of the percentages for the patients seen for Scenario 1 versus Scenario 2 are very close, the same, in fact, for the new standards. This indicates that, for the simulation model and the parameters used, eight residents are too many, and perhaps six should be used instead.

7.3 Cost-to-Patient Throughput Ratio

Finally, we examine a final performance measure, cost-to-patient throughput ratio.

Baseline: -$7295.10 / 17.5833 = -$414.89 per patient
Scenario 1: -$7213.90 / 15.5167 = -$464.91
Scenario 2: -$6916.49 / 15.9167 = -$434.54
Scenario 3: -$7147.79 / 16.0000 = -$446.74

New Standard: -$7553.66 / 16.3333 = -$462.47 per patient

Scenario 1: -$7520.79 / 16.9000 = -$445.02
Scenario 2: -$7253.09 / 16.9000 = -$429.18
Scenario 3: -$7534.91 / 16.2167 = -$464.64

Here, it is evident that for the baseline models, Scenario 2 yields the smallest loss while in the new standard models, Scenario 2 also returns the smallest loss. This confirms the conclusions drawn earlier, that eight residents are too many for the parameters used in the simulation model.

7.4 Patient Throughput Level

The seemingly low number of patients seen can be addressed as follows. The simulation cannot produce more patients than the number of patients that go into the system, i.e., the arrival rate data determines the number of patients seen. Therefore, in order to have a high number of patients leave the system, the arrival rate needs to be greater than what it is in the base model. The health care management group believes that it is possible to produce 160 patients each day, based on standards developed by the benchmarking team. Since the current arrival rate of the patients is well below 160, it is impossible to achieve that level of output. However, suppose that it is possible to input 160 patients into the system. Can the clinic process 160 patients in one business day? From the previous analysis, it is believed that 160 patients are not possible. To demonstrate this, a simple experiment has been set up using the proposed standards shown above. The input parameters are now distribution-free standards determined by the benchmarking team that they believe is a level the staff can achieve. The revised standards are listed below.

Registration		2 minutes
Vitals/Weight		3
R1 Initial Consultation	Return	13
	New	18
R2 Initial Consultation	Return	10
	New	18
R1 Preceptor	Return	6
	New	7
R2 Preceptor	Return	6
	New	6
R1 Final Diagnosis	Return	2
	New	3
R2 Final Diagnosis	Return	2
	New	4
Lab Test		6
Receptionist		10

Each experiment consisted of 600 replications, and each replication ran for a period of nine hours. Nine hours (from 8 AM to 5 PM) allows the clinic staff approximately one hour to clear the system of patients after the last arrival. After the first experiment, the receptionist station is determined to be the bottleneck station with only a mean of 94 patients exiting the system after each replication. Therefore, the standard for the receptionist was lowered to 8 minutes. Once again, the receptionist station is still the bottleneck; however, the time in queue at the station was reduced by over 43 minutes. The mean number of patients leaving the system was increased to 118. Thus, the standard was lowered once again, to 7 minutes. In this case, the R1 queue is now the bottleneck, and the number of patients exiting increases again to 132 patients.

Therefore, the R1 standards are lowered to 12, 11, and 10 minutes, with virtually no change in the number of patients exited, 133. At this point, with the bottleneck station still the R1's, the receptionist standard was lowered to five minutes, and the number of patients exited increased to 140. This indicates that even though the R1 station is a bottleneck in terms of queue length, there are other factors that are causing the number of patients to remain constant. Thus, the preceptor standard was lowered to 5 minutes, and the number of patients exiting rose to 149. The final values for the input parameters that changed are:

R1 Initial Consultation	Return	10
R1 Preceptor	Return	5
	New	5
R2 Preceptor	Return	5
	New	5
Receptionist		5

These standards have reached the limit in terms of being achievable; anything lower would be very difficult to attain. Since the final number of patients leaving the clinic is 149, it can be concluded that most likely, 160 patients exiting the clinic cannot be achieved. Even 149 would be very difficult to achieve, since these input parameters are discrete time values and do not take into account the inherent variability of the system.

CONCLUSION

The health care management group was interested in two main areas of performance for their clinic:

- Maximizing patient throughput, and
- Minimizing patient flowtime.

Based on the initial three scenarios, simulation has shown that of the three scenarios, the second scenario utilizing:

- 6 R2's, and
- 2 medical assistants,

is an optimal staff size. Using performance measures called Vexpenses and Revenues, which are the variable expenses and the patient revenues, respectively, this conclusion has been further solidified as the best among the alternatives.

Proposed changes to the clinic include revising current standards to those proposed in the *Analysis* section. However, it has been shown that based on those standards, the desired goal for number of patients through the system, 160, would be very difficult if not near impossible to achieve. Therefore, a proposed solution would be to lower the expected number of patients at the current level of staffing. It is believed that changing the number of support staff would enable the clinic to reach their goal of 160 patients, and an experimental design investigating the staff size will be the subject of future work. Otherwise, efforts to bring in at least 160 patients per day are required to reach levels anywhere near 160 patients seen.

There is no mention of implementation results, since this analysis still needs to be presented to the doctors and staff at the clinic. Pending the outcome of the presentation and how well this analysis is received by the personnel, implementation will follow shortly thereafter. A favorable outcome may be more likely if the doctors can understand that simulation is a credible tool that can help streamline an operation that is currently running sub-optimal. As the inclusion of Vexpenses and Revenues would indicate, cost is the driving force behind a system, and its incorporation as a performance measure should help both the health care management group and the doctors and staff better understand how changes can be made to the clinic. Simulation allows the user to make changes to the clinic that would be otherwise impossible to implement.

REFERENCES

Banks, J., J. S. Carson, II; and B. L. Nelson. 1996. *Discrete-Event System Simulation*, 2nd Edition, Prentice-Hall, Upper Saddle River, New Jersey.

Kelton, W. D., R. P. Sadowski, and D. A. Sadowski. 1998. *Simulation with Arena*, McGraw-Hill, Inc., New York.

Law, A. M., and W. D. Kelton. 1991. *Simulation Modeling and Analysis*, 2nd Edition, McGraw-Hill, Inc., New York.

Lowery, J. C. 1998. Getting Started in Simulation in Healthcare. In *Proceedings of the 1998 Winter Simulation Conference*, ed. D.J. Medeiros, E.F. Watson, J.S. Carson, and M.S. Manivannan, Institute of Electrical and Electronic Engineers, Piscataway, New Jersery

Nelson, B. L. 1995. *Stochastic Modeling: Analysis and Simulation*, McGraw-Hill, Inc., New York.

AUTHOR BIOGRAPHIES

MARK L. WENG is a M.S. candidate at the University of Cincinnati. He holds a B.S. in Industrial Engineering from Northwestern University. His research interests include the statistical aspects of simulation modeling, stochastic models, and applications of simulation, particularly in the health care field.

ALI A. HOUSHMAND is an Associate Professor of Industrial Engineering at the University of Cincinnati. He has a B.S. and M.S. in Mathematics from the University of Essex, UK, and M.S. and Ph.D. degrees in Industrial Engineering from the University of Michigan. His research and teaching interests are in applied statistics and in quality control, design, and management.