SIMULATING A PROPOSED INFORMATION DELIVERY SYSTEM

Eugenio M. Alvarez

Nine West Group, Inc.
9 West Broad Street
Stamford, Connecticut 06902, U.S.A

Paul Mitchell
Raymond Telerole

New York Power Authority
123 Main Street
White Plains, New York 10601, U.S.A

Walter Callahan

New York Power Authority
Indian Point Nuclear Power Plant Unit 3
Buchanon, New York 10511, U.S.A

Peter Hoefer

Lubin School of Business
Pace University
Pace Plaza
New York, New York 10038, U.S.A

ABSTRACT

This paper describes the results of a simulation analysis of hardware designs for a proposed data delivery system. The system is to serve clients of a large accounting firm. Once the required hardware components were identified, and their system parameters and response times estimated, the design of such a system required selection of components in a way to optimize specific performance characteristics. Simulation was used to understand the performance of the system, and to help with decisions on the components to be used.

1 INTRODUCTION

This story begins with one of us being involved with designing a system to give customers an interactive ability to retrieve information from a large accounting firm. In the next section we describe the design of this "InfoDelivery System," a data storage and acquisition system, intended to give users current information on time and expense records. The system design enables the user to access the information through on-line reports that are dynamically linked. The user starts with a report and "drills down" to the linked reports the user wishes to access. It is expected that this system will eventually replace numerous reporting systems currently in use.

Since speed and system reliability are vitally important to users of this system, basing the system on single components may not be acceptable. However, some duplicate components are expensive. This study will not describe the InfoDelivery System in detail, but will provide an analysis of the critical component mix required to access the system. In this context, the relevant variables are speed of access and cost.

2 SYSTEM DESIGN OVERVIEW

The InfoDelivery System contains three subsystems that link together to create the application as a whole: the client computer (PC), the network, and the mainframe database.

The speed of data retrieval is vitally important to the application. During the analysis of this system, customer "throughput time" was closely monitored. Early in the study, we learned that acceptable time was an average 20 seconds or less per customer. Thus, components had to be selected in a way that allows average throughput time of 20 seconds or less, while containing cost. The discussion in this paper mainly focuses on throughput time; the information provided to management enabled them to decide which design they could accept based upon cost and performance.

The critical hardware requirements and connections are described in Figure 1 (see next page). There are three components of interest associated with six variables (the variable names are capitalized in parentheses):

1. The PC client, which consists of
 - A. A security module(s) (SECURITY)
 - B. A data access layer(s) (DAL)
 - C. A swap pool(s) (SWAP)
 - D. An error handler(s) (ERRH)
2. The network (an ODBC communication layer(s))
3. The DB2 mainframe database(s).

We treated each of the six variables (SECURITY, DAL, SWAP, ERRH, ODBC and DB2) as discrete components that could be connected in parallel. Thus, if a queue formed for DB2, a second DB2 component could be added to reduce that queue. In effect, adding components was modeled as adding parallel identical servers. A technical session with

Figure 1: Hardware Requirements and Connections

management, not involving costs, led to the determination of the feasible values of the variables, as follows:

The value of SECURITY was determined to be feasible between 1 and 3;
 the value of DAL was determined to be feasible between 1 and 2;
 the value of SWAP was determined to be feasible between 1 and 2;
 the value of ERRH was determined to be feasible between 1 and 2;
 the value of ODBC was determined to be feasible between 2 and 3;
 the value of DB2 was determined to be feasible between 1 and 3.

The cost per component varied greatly. The estimates are as follows:

Variable	Cost per Unit
SECURITY	\$500.00
ODBC	\$3,000.00
DB2	\$50,000.00
DAL	\$100.00
SWAP	\$75.00
ERRH	\$10.00

Clearly, the large variability of per unit cost affected the decision process.

Other parameters that are important for this paper were either estimated, or known to be, as follows:

Maximum number of users: 5,000.

Three periods of system load over a cycle (10 business days):

- Peak period - 5,000 users within a 16 hour (2 business day) period
- Normal period - 2,500 users within a 48 hour (6 business day) period
- Off period - 500 users within a 16 hour period (2 business day) period.

In an initial analysis, what is obvious was confirmed; the peak period use dominates the decision about component mix. This paper therefore focuses on peak period use, although the entire study has information on all periods.

The exponential distribution was used to model interarrival times for requests for service, as determined by the peak period requirements. Service times were modeled by the uniform distribution, with estimated mean and half-widths (in seconds) as follows:

SECURITY	first time	1, .25
	second time	.5, .25

DAL	1.5, .75
SWAP	1, .75
ERRH	.5, .25
ODBC	3, 1.5
DB2	4, 2.5

Experiments using these parameters with certain different distributions did not show the decision options were very sensitive to those changes (although the throughput time averages increased with distributions containing more variation). Thus, the report contains performance based upon the uniform distribution, without compromising the decision making process.

3 THE SIMULATION MODEL AND EXPERIMENT

The systems described above can be easily modeled in any modern simulation language. We chose GPSS/H (Wolverine (1989)). We studied a base model for sensitivity to various model changes, and to determine the approximate number of replications necessary to have a reasonable error term in confidence intervals for throughput time (Schriber (1990)). Experts observed the simulated model performance, and suggested changes so the simulated system behaved like what they believed was an actual working system. With four of the six variables having two levels of treatment (SWAP, DAL, ERRH and ODBC), and the other two variables having three levels of treatment (SECURITY and DB2), a complete analysis consisted of evaluating the performance of 144 models (Box (1978), Madu (1993)). We used 25 replications for each model, looking at the peak performance period only. Five recommended designs, with mean throughput time and cost, are listed in Table 1. Note

that there are statistical “ties” between some configurations (no statistical significance was detected). Using ANOVA, a significant difference was detected among the five configurations in Table 1. The last configuration is significantly slower (and cheaper) than the other configurations listed in Table 1; the first configuration is significantly faster (and more expensive).

4 CONCLUSION: DISCUSSION OF THE RESULTS

There were many model designs that met the less than 20 second average throughput requirement in a statistically significant sense. Some of those designs came close to the theoretical minimum of just more than 13.5 seconds. Table 1 lists five designs which distinguish themselves from other designs because of sample throughput ranking and cost.

The results show that for “optimum” performance, at least three thousand dollars must be added to the least cost design that meets the required performance criterion. The design with the “best” simulated average time, listed at the top of Table 1, is extremely costly because of the third DB2 component. Even though there is a significant difference among the performances of that design and other less costly designs (such as the middle three in Table 1), we do not recommend it because of the cost. There is also a significant difference among the throughput times for the last design listed on Table 1 and the four above it. We nevertheless suggest any of the last four, with a stipulation that other information on reliability and other practical measures be taken into consideration before selecting a particular model. For example, if there is suspicion that the error handlers will not be very reliable, given their cheap cost, we recommend purchasing more of them than may be needed.

ACKNOWLEDGEMENTS

We thank Dr. Christian Madu, Professor of Management Science and Research Professor at the Lubin School of Business, Pace University, for the suggestions he made to improving the analysis in this paper.

REFERENCES

- Box, George E. P., William G Hunter and J. Stuart Hunter. 1978. *Statistics for Experimenters*. New York, NY. Wiley-Interscience.
- Madu, Christian N., and Chu-Hua Kuei. 1993. *Experimental Statistical Design and Analysis in Simulation Modeling*. Westport, CT. Quantum.
- Schriber, Thomas J. 1990. *Introduction to Simulation Using GPSS/H*. New York, NY. Wiley and Sons.
- Wolverine Software Corporation. 1989. *GPSS/H Reference Manual*. 3rd Edition. Annandale, VA.

AUTHOR BIOGRAPHIES

EUGENIO M. ALVAREZ is the Data Administrator for Nine West Group Inc. He received a B.A. degree in Computer Information Science from Iona College, and an M.S. degree in Management Science from Pace University. He has previously worked for AT&T, the New York Power Authority, and most recently, KPMG Peat Marwick.

WALTER M. CALLAHAN is the Senior Accountant at the Indian Point #3 Nuclear Power Plant, a subsidiary of the New York Power Authority. He holds a B.A. degree in Social Sciences from SUNY Binghamton, and an M.S. degree in Management Science from Pace University.

Table 1: Selected Designs: Performance and Cost

Design*	Estimated Cost	Mean Throughput: Peak Period	Comments
(3,3,3,2,2,1)	\$160,860	13.58 secs	Best simulated average time
(3,3,2,2,2,1)	\$110,860	13.84 secs	
(2,3,2,2,2,1)	\$110,360	13.86 secs	
(1,3,2,2,2,1)	\$109,860	13.89 secs	
(1,2,2,1,1,1)	\$106,685	15.77 secs	Least cost design with speed < 20 secs

*The six-tuple (A,B,C,D,E,F) indicates the number of components for A=SECURITY, B=ODBC, C=DB2, D=DAL, E=SWAP, F=ERRH.

PAUL R. MITCHELL is a Senior Engineer in the Performance Engineering group of the New York Power Authority. Previously, he was a senior engineer with Niagara Mohawk Power Corporation and a principal engineer with General Physics Corp. He received a B.S. in Mechanical Engineering from the University of Arizona in 1970, and an M.S. in Management Science from Pace University in 1996.

RAY TELAROLE is Manager of Technical Training for the New York Power Authority Power Generation Division. He received a B.A. degree in Humanities from St. Thomas Seminary College, and an M.S. degree in Management Science from Pace University. He has previously worked for the Utah Power and Light Company, and most recently General Physics Corporation.

PETER HOEFER is Professor of Management Science and Associate Dean and Director of Graduate Programs at Pace University's Lubin School of Business. He holds a Ph.D. degree in Mathematics from the City University of New York. In addition to administration, his interests include risk analysis and simulation.