USE OF SIMULATION TO TEST THE VALIDITY AND SENSITIVITY OF AN ANALYTICAL MODEL

PROSPER M. BERNARD

The Provincial Bank of Canada

ABSTRACT

The purpose of the research reported here was to test formally the validity of some assumptions made in solving models by analytical techniques and to test the sensitivity of the system to the arrival distribution. The presentation refers to a simplified real-time model. The analytical solution in the literature is to derive shown results for each stage of the system and add them up to obtain the behavior of the system. Assumptions must then be made at each stage. Simulation has been used to solve the same system in terms of the same measure of efficiency, i. e. the response time. However, the system is solved as a whole, the output from one stage becoming the input to the second stage. Confidence limits have been obtained for the response time in order to test the results obtained from analytical techniques. Simulation has also been used to test the sensitivity of the system to a change in the arrival distribution. Using analysis of variance, the effect of the arrival pattern and of the interaction is determined.

Introduction

The purpose of this presentation is to show how one can use simulation to verify formally some assumptions made in solving a model by analytical techniques and to test the sensitivity of that model to a departure from the classical

"Poisson" arrival assumption. The method used in this paper is to describe the model, which has been published and to discuss an analytical solution. As will be seen in detail later, the procedure for solving this multi-stage model of queues in series is first to solve each stage and then to add up the results. This implies making assumptions at each stage, and this procedure may introduce errors. For this reason, the simulation will be used to study the system as a whole without having to make intermediate assumptions. Independent simulation experiments will be made in order to obtain confidence limits on the mean response time.

Most queuing models assume a Poisson arrival distribution and no solution is offered for a departure from the "Poisson" assumption. It might be easier to determine in advance the effect of the arrival distribution than to make field studies to make sure that it is in fact Poisson. Even if one could determine the exact shape of the distribution and it were found not to be the Poisson distribution, there is no solution available. For this reason the author will test the sensitivity of the model by repeating simulation experiments with distributions as far apart as Poisson, Normal, and Uniform.

Model to be tested

In this real-time inquiry model there are a finite number of customers linked to a central processor via common carrier lines and a terminal. A graphical representation of the model is given in figure 1.

The following assumptions have been made:

1. The arrival distribution at each terminal is Poisson with mean λ/m where m is the number of terminals.
2. The service time of the central processor is distributed according to an Erlang-2 distribution.
3. The service time for each terminal, i. e. the key-in time and print-out, are uniformly distributed.
4. The queue discipline is First-in-First-out.

The real-time inquiry model considered in this presentation could be seen as a machine interference model where there are external arrivals to each machine or, in this case, to each terminal. Customers arrive at each of the m terminals randomly with mean arrival rate λ/m. The total response time consists of the waiting-time for the terminal, the service-time of the terminal (key-in), the waiting for the CPU, the service-time in CPU and the service-time of the terminal (print-out). The system is consi-

dered busy during the whole period and can accept a new customer only after the terminal processing of the previous customer.
The reader will see the similarity between this model and the classical "machine interference" model. The whole model can be seen as a multi-stage model where the middle stage is a machine interference model. This latter model has been well described and analyzed in the literature.
In the machine interference model there are a finite number of machines or sources assigned to one serviceman or service station. The machine is either "up" or "down". When the machine goes down it joins the queue for service. The machine gets immediate service or waits for service depending on the availability of the single repairman.

In the machine interference model the following assumptions are usually made:

1. The service time is exponentially distributed with mean $\bar{T}s$.
2. The "up" time for each machine is exponentially distributed with mean time $\bar{T}a$.

These assumptions are sometimes referred to as the worst-case conditions. The ratio of the two times is defined as the "service ratio"

where

$$Z = \frac{\bar{T}a}{Ts}.$$

The "Server Utilization" denoted as $r_{m(z)}$ can be obtained for each number of machines m and for each value of z as follows. Let the probability P_0 represent the fraction of time when there are zero machines in the service queue, and the serviceman is idle. Thus, $1 - P_0$ may be interpreted as the fraction of time the serviceman is busy, that is:

$$\text{Server Utilization} = 1-P_0 = 1 - \frac{e^{-z}\frac{z^m}{m!}}{e^{-z}\sum_{j=0}^{m}\frac{z^j}{j!}} = r_{m(z)}.$$

Analytical solution

The model has been given not a rigorous solution, but an approximation, where the results of all stages are added up to make the total response time. The following solution has been proposed [1].
From the user's point of view, the system is in use when he starts keying the inquiry so that it could be considered as a service-station or a "black box" in which service time Tp is:

$$Tp = Tu + Tw + Ts + To \qquad 1$$

where

Tu = Time to key-in and transmit the message

Tw = Waiting time to access CPU

Ts = Time for service in the CPU

To = Time to transmit and print the result

Tp = Total service time

Tq = Total response time = waiting for the system + Tp.

The model is represented in figure 1.

Since the overall system can be considered to be a single service station where the expected service-time is $\overline{T}p$, the system utilization by a user can be defined as:

$$\rho = (\lambda/m) \; / \; (1/\overline{T}p) = \frac{\lambda}{m} \cdot \overline{T}p \qquad 2$$

The response time for a single server model with random arrivals and an arbitrary service distribution has been obtained by Pollaczek and simplified by Khintchine[2]. The general formula known as Pollaczek-Khintchine uses only the first two moments of the service distribution and can be transformed by algebraic modifications to:

$$Tq = \frac{Tp}{1-\rho} \left[1 - \rho/2 \; (1 - \frac{\sigma^2}{Tp^2})\right] \qquad 3$$

The total response-time has been obtained in terms of Tp, the service-time when the overall system is assumed to be a service-station. The only further information necessary to obtain the expected value of Tp is the expected waiting time to the CPU itself.

By applying the machine interference results for the middle subsystem where the machines or terminals are queuing the CPU, one can obtain the remaining values. The server's (CPU) utilization R(m)Z can be found from the graph in figure 2. However, the service ratio z is not known, since it depends on the external arrival rate. This ratio can be determined in the following manner. Since all customers arriving at the various machines must eventually go through the service queue, the utilization of the serviceman can be calculated, independently from z, to be

$$r_m(z) = \lambda . \overline{T}_s \qquad 4$$

Z $\overline{T}s = \overline{T}a$ and $r_m(z) = \lambda \overline{T}s$ can be substituted in the general machine interference formula

E(time between breakdowns) = $\frac{mts}{r_m(z)}$ - ta.

The following results are thus obtained.

$$Tw + Ts = M/\lambda - Z \; Ts \quad \text{if} \quad \frac{\overline{T}w}{Ts} > 1. \qquad 5$$

However, for $\frac{\overline{T}w}{Ts} < 1$, the simple queuing time formula may be used as a good approximation.

$$Tw + Ts = \frac{Ts}{1 - \frac{(m-1)}{m} \lambda \; Ts} \quad \text{if} \quad \frac{Tw}{Ts} < 1 \qquad 6$$

where Z is the service ratio in the machine interference model.

Example:

The behavior of the model can be shown better in terms of an example. In this example there are 20 terminals connected to the CPU. The key-in time is uniformly distributed between

5 and 15 and the print-out of a message is also uniformly distributed between 2 and 7. The computer processing time is assumed to be an Erlang-2 distribution with a mean of 2 seconds. Although the CPU processing time is not exponentially distributed, the machine interference formulae are used as an approximation.

In order to obtain numerical values for the expected response-time and for the other components one may proceed as follows:

CPU utilization is determined from equation

$$r_{20}(z) = \lambda \frac{\text{Inquiries}}{\text{sec}} \times \frac{2 \text{ sec}}{\text{inquiry}}$$

From this equation the service ratio z is determined and can then be substituted in equations 5 and 6.

$$\overline{T}_w + \overline{T}_s = \frac{2}{1-0.95 \times 2 \times \lambda} \quad \text{if } \frac{Tw}{Ts} < 1$$

or

$$\frac{20}{\lambda} - 2z, \text{ if } \frac{Tw}{Ts} > 1.$$

The overall inquiry service time found from equation 1 is then given as

$$Tp = 10 + Tw + Ts + 5.$$

The inquiry service variance is likewise given as the sum of variances

$$\sigma_p^2 = \frac{(15-5)^2}{} + Tw^2 + \frac{2^2}{2} + \frac{(7-3)^2}{12} = 11.7 + \overline{T}w^2$$

where σw^2 is approximated by Tw^2.

Terminal utilization is given from equation 2.

$$\rho \frac{\lambda}{20} = Tp.$$

Finally, the inquiry response time is determined from the queuing time formula of equation 3. The whole set of calculations is summarized in Table 1.

With 20 terminals the system can accept 0.4 inquiries per second or 1.2 inquiries per minute per terminal. The response time is 30 seconds with CPU utilization of 80%. Beyond this point the response time increases at an increasing rate as shown in figure 5.

The reader will notice that assumptions have been made above concerning the arrival distribution at the second stage of the system. Burke has shown that the output of one queue with Poisson input is also Poisson[3]. However, the general procedure for obtaining the distribution of the output of queues in parallel and series is difficult to obtain analytically. In this analytical model, assumptions have to be made at each stage. It is assumed that the input to the second stage is Poisson as weel as the input to the system.

In solving the model, assumptions were made and approximations were used. The model was analytically solved by studying each of the subsystems and adding up the results. Intermediate assumptions were made concerning each subsystem. The middle stage was approximated by the machine interference model where the "down" times are exponentially distributed. Even in that middle stage, the usual queuing theory formulae for arrivals from an infinite population were used where the ratio of the waiting time over the service-time at the CPU was small.

Like any abstraction, the model represents only some aspects of the real world. It is important, therefore, to know the effect of assuming or neglecting certain parameters. The model is based on the "Poisson" arrival assumption and no solution is given for a non-Poisson arrival. It may be asked what would happen if the arrival distribution was not really Poisson. To explore this, the two following hypotheses will be tested.

Hypotheses

A. The expected total response time is properly obtained from the analytical solution.

B. The arrival distribution has no effect on the expected total response time of the simulation.

Simulation results

Simulation runs using GPSS were made to test the above hypotheses. The simulation runs are made for the same values as in the example given above. Th program generates Poisson arrivals randomly distributed to any of the 20 terminals. A queue is formed in front of each terminal. The units go through the stages of the model illustrated in figure 1. In the simulation, there is no need to make assumptions at each stage since the program will take as input in stage 2 the output of stage 1. Measurements are made only at the end, for the customer is interested in knowing the response time of the overall system rather than the waiting at each stage.

The First Hypothesis. The system was studied at eight different arrival rates corresponding to the rates shown in table 1. In order to arrive at a confidence interval at each point of interest one must have a sample of independent observations. However, data generated by simulation are autocorrelated. If we assume the absence of autocorrelation we may underestimate the variances or we may take a too small sample. The variance of autocorrelated data is not related to the population by the simple expression

$$\sigma^2_{\bar{x}} = \sigma^2/n$$

but by

$$\sigma^2_{\bar{x}} = \sigma^2/n + k$$

where k is a positive number. In order to avoid the problem of autocorrelation, 12 independent runs were made at each of the 8 arrival rates being studied. The mean of each run was used. In each run a transient period of 50 arriving units was discarded and 100 steady state units were recorded. By making exploratory runs it was found that the steady-state was reached well before 50 units had arrived. The mean of the response time for each is approximately normally distributed and a confidence interval can be calculated.

It is not possible to show here the results of so many runs. However, the expected response time for each run was recorded and entered in the first column of table 2.

The output of the simulation runs are compared to the analytical results in table 3. Although our interest here is in the expected total response time, table 2 also shows the utilization values, the waiting times and the length of the queue. However, the response time has not been obtained by adding up the different items but by measuring the difference of time between the arrival and the departure of a unit. The other values are presented to help identify the area of great differences, as will be discussed later. Both response times are graphed in figure 6. It will be noticed that there is a major difference in the results when the arrival rate approaches 1.0 unit per second or when the CPU utilization approaches 1.

The analytical results imply that the mean response time curve shown on figure 1 approaches asymptotically a vertical line at an arrival rate below 0.5 inquiry per second. The simulation results show that the expected response time curve approaches asymptotically a vertical line at 0.75 inquiry per second.

It should be noted here that the results of the simulation model were not obtained starting from empty system. This would have produced results further away from the analytical results. Instead, the model was run until steady state had been reached. Only at that point were the statistics accumulated.

As one will recall, the analytical solution was obtained by adding the waiting time and the service time at each stage. One can see, by looking at table 3, that the element that varies the most between the analytical and the simulation results is the waiting time at the CPU itself. This corresponds to the waiting time at the stage that was approximated by the machine interference model. It is interesting to note that in the area where the ratio of the average waiting time to the CPU over the average service time was less than one (i.e.) $\frac{\bar{T}w}{Ts} < 1$), the results of the analytical solution and of the simulation are statistically the same. In this region the machine interference model was not used.

Confidence intervals for the total response time were calculated using Student t Statistics for the 8 arrival rates of interest at 99% level and 95% level. The results are summarized in table 4 and plotted in figure 4. It can easily be seen that for an arrival rate of over .45 inquiry per second, the analytical results are well outside the 99% confidence interval.

It is easy to conclude that the results fail to confirm the first hypothesis. The response time does not increase as fast as suggested by the analytical techniques. The reason for this may be that the input to the second stage of the model, i.e. the CPU, is not exponentially distributed, as has been assumed to fit the machine interference model. The later model assumed exponentially distributed "down" times on each of the 20 terminals. However, since a queue is formed in front of each terminal followed by a uniform service time, the equivalent of the "down" times are not necessarily exponentially distributed. The reader will notice that in the simulation runs there is no need to make assumptions at the second and each subsequent stage as in the analytical solution.

Testing the Second Hypothesis. An important factor often mentioned in the literature is that if one wants to use the arrival rate, a special study should be undertaken to make sure that it is Poisson. In this presentation the author uses another approach, i.e. sensitivity analysis of the assumptions. One would ask What if the assumptions are not true? Should we investigate the real arrival pattern?

It is obvious that there is no need to investigate the exact value of a factor if this factor has no effect on the system. The author's approach is therefore, to determine in advance the effect on the system of the mean of the arrival distribution. In order to do this, the author repeated the experiments described above for two other distinctive arrival patterns, i.e. normal and uniform arrival. The means of the three distributions are the same. The uniform distribution varies from 0 to $2\overline{X}$ where $\overline{X}$ is the mean. The standard deviation used for the normal distribution is in this case $\overline{X}/5$. In each case, as before, 12 runs of 100 steady state arrival were made for each of the 8 levels of interest. The results obtained are summarized in table 2. There are 288 runs of 100 observations, i.e. 28,800 observations in all, not including the transient period. We also have enough information to test the interaction since each cell in the table has 12 observations.

An analysis of variance was made to determine the effect of the treatment (arrival distribution) on the response time. The arrival distributions and the arrival means are both fixed factors. There are 3 levels for the first and 8 levels for the second factor and 12 independent observations in each cell as shown in table 2.

The arrival means have an obvious effect on the mean response time. However, here one is interested in testing the effect of the arrival distribution and the interaction effect. The results of the analysis of variance are shown in table 5. The surprising conclusion is that the arrival distribution has no significant effect on the total response time (even at a low 73% confidence level) and that the interaction, i. e. the combined effect of arrival rate and arrival pattern, is almost non-existent. We do not reject the second hypothesis.

This conclusion implies that there is no need to search for the true arrival pattern since it has little effect. Even if the arrival to the system exhibits a departure from the Poisson assumption the expected response time seems not to be affected.

Conclusion

In this paper the author has reviewed and tested by simulation a simplified real-time model. He has shown that the system does not behave as prescribed by the analytical solution and has demonstrated the original conclusion that the arrival distribution has little effect on the expected total response time. However, a more extensive study made by the author on many other models showed that many models are sensitive to a departure from the Poisson arrival.[5] Simulation has been used successfully to arrive at this conclusion because there is no need for assumptions to be made at each stage as in the analytical solution. Furthermore, a simulation model, once running, can easily be changed to analyse the assumptions at the first stage. It is obviously less expensive to run sensitivity analysis on the assumptions than to field-test them. Sensitivity analysis could be performed on any of the other parameters of the model. However, simulation experiments being very expensive, the author selected to allocate limited resources to the analysis of the classical queuing theory assumption, i. e. the Poisson arrival distribution, in a simple model.

Biography

Prosper M. Bernard holds a B.A. from University of Montreal, a B. Sc. from McGill University, and an MBA from St. John's University New York. The degree of Doctor of Philosophy is expected soon from the City University of New York. He is also a CDP holder. Mr. Bernard held teaching positions both in Quebec and in New York. He also has a long experience in the information processing field as a programmer, systems analyst, manager, and management consultant. He is presently with the Provincial Bank of Canada as an internal consultant.

Acknowledgment

The author would like to thank Dr. Lloyd Rosenberg of the City University of New York for his helpful guidance during his research. The author is also grateful to Maritime Overseas Corporation and Israel Discount Bank in New York for the free use of their computer in doing this research. The Provincial Bank of Canada has also being instrumental in making this presentation possible.

References

1. Analysis of Some Queuing Models in Real-Time Systems, (New York: IBM, Ref.: F20-0007-1, 1966), p. 44-46.
2. T. L. Saaty, Elements of Queuing Theory, New York: McGraw Hill, . 1961) p. 40.
3. P. J. Burke, "The Output of a Queuing System", Operations Research, Vol. 4, December 1956, pp. 699-704.
4. Prosper M. Bernard, "Use of Computer Simulation to Test the Validity and Sensitivity of Real-Time and Time-Sharing Queuing Models", PhD Thesis, City University of New York, 1973.

<u>Table 1</u>

<u>Example of the results obtained from the analytical solution</u>

CPU Utilization R_{20}	Arrival rate	Service ratio Z	Waiting time CPU $\overline{T}_w$	Terminal utilization	Total time in system $\overline{T}_q$
0.1	0.05	200	0.21	0.04	17.5
0.2	0.10	98	0.47	0.09	18.5
0.4	0.20	48	41.20	0.18	20.2
0.6	0.30	31	22.65	0.30	24.2
0.8	0.40	21	64.80	0.44	31.2
0.9	0.45	17	67.30	0.55	41.0
0.98	0.49	13	212.5	0.72	75.0

TABLE 2

SIMULATION

EXPECTED RESPONSE TIME

MEAN ARRIVAL RATE SEC.	ARRIVAL DISTRIBUTION: POISSON	UNIFORM	NORMAL
0.10	16.7 17.3 18.4 17.6	17.9 18.8 18.2 16.1	17.5 17.2 17.7 17.9
	17.5 18.0 18.5 18.3	18.7 17.6 17.3 17.9	16.6 17.4 17.3 17.3
	18.3 17.6 19.2 17.8	18.3 17.9 17.0 18.3	18.2 18.4 17.9 17.7
0.20	25.9 21.6 19.4 18.6	18.8 19.7 19.4 18.8	19.1 17.9 18.6 17.3
	19.7 19.8 19.8 22.4	20.6 20.8 18.3 20.3	18.9 18.9 18.4 19.0
	20.9 20.2 18.7 19.6	18.0 20.1 19.0 21.3	19.5 18.6 18.9 19.7
0.30	27.9 23.3 25.8 21.0	22.4 21.6 19.8 20.5	22.4 22.9 20.4 20.3
	20.4 25.1 22.5 27.1	18.9 27.3 22.6 22.4	20.8 20.5 21.3 22.3
	23.1 20.2 21.0 22.4	23.3 20.8 19.9 21.6	21.2 23.6 21.8 22.8
0.40	44.1 32.4 27.2 26.4	33.5 24.5 20.9 26.2	27.1 24.3 22.1 23.1
	28.2 21.0 23.1 23.1	26.5 36.2 24.5 27.2	23.9 27.2 26.0 24.6
	28.1 29.6 26.1 33.3	25.4 26.4 32.7 37.9	21.5 27.3 26.7 27.6
0.45	41.6 27.7 32.1 43.5	32.1 29.7 26.6 32.5	32.7 26.6 26.9 29.0
	31.2 27.1 37.8 38.4	29.1 27.3 26.4 25.0	30.0 31.8 23.7 25.6
	31.3 24.4 27.7 29.3	50.0 32.4 29.8 29.6	26.1 30.2 29.6 34.2
0.50	31.9 33.1 57.3 28.9	38.3 29.3 46.4 37.4	40.7 32.5 32.2 29.0
	39.1 37.3 30.9 41.8	43.8 39.3 33.7 44.8	32.3 46.9 41.0 28.2
	59.6 39.7 59.1 70.2	58.6 40.6 32.4 33.7	31.5 39.6 34.3 39.4
0.555	47.5 39.8 35.2 61.7	55.3 33.1 41.1 43.2	34.4 43.7 34.7 60.0
	40.0 29.4 36.2 25.4	61.5 46.8 40.2 54.0	39.1 60.6 55.2 41.0
	90.6 54.7 37.4 56.5	51.2 48.3 45.6 52.2	39.5 45.3 59.7 62.1
0.566	65.6 106.0 43.7 49.2	57.3 65.0 83.4 73.9	69.5 45.5 76.6 83.8
	95.8 91.1 60.2 57.2	64.2 72.5 56.2 80.8	76.2 70.0 67.8 69.4
	71.2 74.9 63.0 71.5	77.8 61.4 45.6 56.7	64.1 78.2 51.0 71.3

NOTE: EACH VALUE SHOWN IN THIS TABLE IS THE AVERAGE OF OA RUN MADE OF 100 NON TRANSIENT OBSERVATIONS.
THERE ARE 283 RUNS WITH DIFFERENT RANDOM NUMBER SEQUENCES.

TABLE 3

A N A L Y T I C A L   R E S U L T S   V S   S I M U L A T I O N   R E S U L T S

INTER ARRIVAL TIME (SEC) ANAL.	INTER ARRIVAL TIME (SEC) SIM.	ARRIVAL RATE ANAL.	ARRIVAL RATE SIM	UTILIZATION OF CPU ANAL.	UTILIZATION OF CPU SIM.	WAITING FOR CPU ANAL.	WAITING FOR CPU SIMU.	TERMINAL UTILIZ. ANAL.	TERMINAL UTILIZ. SIM.	RESPONSE TIME EXCL. ANAL.	RESPONSE TIME TERM SIM.	RESPONSE TIME INCL. ANAL.	RESPONSE TIME TERM. SIM.	LENGTH OF AVERAGE CPU ANAL.	LENGTH OF AVERAGE CPU SIM.	QUEUE CPU MAXIMUM SIM.
10	10	0.10	0.10	0.2	0.197	0.[illegible]7	0.324	0.09	0.083	17.5	17.25	18.5	18.109	0.047	0.035	2.0
5	5	0.2	0.2	0.4	0.395	1.20	0.928	0.18	0.161	18.2	18.48	20.2	20.012	0.240	0.194	6.0
3.33	3.33	0.3	0.3	0.6	0.606	2.65	1.749	0.30	0.321	19.7	18.50	24.2	22.915	0.795	0.518	7.0
2.5	2.5	0.4	0.4	0.8	0.783	4.80	4.017	0.44	0.475	21.8	22.67	31.2	28.557	1.90	1.771	10.0
2.22	2.22	0.45	0.45	0.9	0.895	7.30	5.563	0.55	0.428	24.3	24.00	41.0	32.724	19.0	2.578	10.0
2.0	2.0	0.5	0.5	1.0	0.961	38.0	8.493	1.0	0.708	55.9	LARGE	LARGE	43.793	LARGE	4.281	13.0
1.8	1.8	0.555	0.555	1.0	0.967	LARGE	4.512	1.0	0.916	LARGE	LARGE	LARGE	46.205	LARGE	4.408	14.0
1.5	1.5	0.666	0.666	1.0	0.999	LARGE	15.601	1.0	0.965	LARGE	LARGE	LARGE	71.016	LARGE	8.516	15.0

ANAL. : RESULTS OF ANALYTICAL SOLUTION
SIM: SIMULATION RESULTS
LARGE: MEANS THAT THE VALUE APPROACHES INFINITY

TABLE 4

C O N F I D E N C E   L I M I T S   F O R   R E S P O N S E   T I M E

ARRIVAL TIME	ARRIVAL RATE	UTILIZATION OF CPU ANAL.	UTILIZATION OF CPU SIMUL. * MEAN OF MEANS	TOTAL TIME SPENT IN SYSTEM = RESPONSE ANAL.	TOTAL TIME SPENT IN SYSTEM = RESPONSE SIMUL * MEAN OF MEANS	STD. ERR. OF MEAN	CONFIDENCE INTERVAL FOR POP. MEAN 95% LEVEL	CONFIDENCE INTERVAL FOR POP. MEAN 99% LEVEL
10	110	.2	0.197	18.5	18.109	0.162857	17.7429 - 18.6042	17.5956 - 18.6042
5	.2	.4	0.395	20.2	20.012	0.326897	19.4145 - 20.8522	19.1213 - 21.1454
3.33	.3	.6	0.606	24.2	22.915	0.614245	21.6160 - 24.3174	21.0651 - 24.8693
2.5	.4	.8	0.783	31.2	28.557	1.76127	24.6770 - 32.4230	23.0974 - 34.0026
2.22	.45	.9	0.895	41.0	32.724	1.79238	28.7170 - 36.5997 ***	27.1094 - 38.2073
2.0	.50	1.0	0.961	LARGE **	43.79	2.52822	28.8989 - 40.0178 ***	26.6314 - 42.2853
1.8	.555	1.0	0.967	LARGE **	46.206	5.13776	34.9023 - 57.4977 ***	30.2943 - 62.1059
1.5	.560	1.0	0.999	LARGE **	71.010	5.41217	58.9572 - 82.7595 ***	54.3031 - 87.6136

* THE VALUE SHOWN IS THE AVERAGE VALUE OF 12 RUNS.
EACH RUN HAS 100 STEADY STATE OBSERVATIONS.

** APPROACHES INFINITY AS UTILIZATION APPROACHES 1.0.
OR AS ARRIVAL RATE APPOACHES 1.0

*** THE RESULTS OF MATHEMATICAL FORMULAE GIVES A RESPONSE
TIME OUTSIDE THE CONFIDENCE INTERVAL.

NOTE : ALL THE VALUES SHOWN ON THIS TABLE ARE THE MEAN
OF ONE SIMULATION RUN OF 100 OBSERVATIONS EXCLUDING
TRANSIENT OBSERVATIONS.
100 OBSERVATIONS FOR 1 ENTRY IN THE TABLE
12 ENTRIES PER CELL
8 LEVELS I.E. 8 ARRIVAL AVERAGES
3 TREATMENTS OR 3 ARRIVAL PATTERN
24 CELLS
288 ENTRIES
28,800 OBSERVATIONS

TABLE 5

EFFECT OF THE ARRIVAL PATTERN ON THE RESPONSE TIME IN A REAL TIME MODEL

ANOV-/-UNWEIGHTED SOLUTION

SOURCE	SS	DF	MS	F
A	74915.250	7.	10702.176	191.928
S(A)	4907.000	88.	55.761	
B	257.766	2.	128.883	1.901
AB	495.891	14.	35.421	0.523
BS(A)	11931.188	176.	67.791	

GRAND MEAN = 34.192

CELL MEANS--ROWS = LEVELS OF A, COLUMNS = LEVELS OF B

	1	2	3
1	18.100	18.000	17.592
2	20.050	19.633	18.733
3	22.967	21.758	21.692
4	28.550	28.483	25.117
5	32.675	30.375	26.933
6	43.775	40.025	35.617
7	46.200	47.7[illegible]8	47.942
8	70.859	66.233	69.100

NOTE: A = ARRIVAL MEANS ; B = ARRIVAL PATTERN
A AND B ARE FIXED FACTORS ; A IS MAIN EFFECT AND IS CONFOUNDED WITH PLOT

Figure 1

A simple real-time model

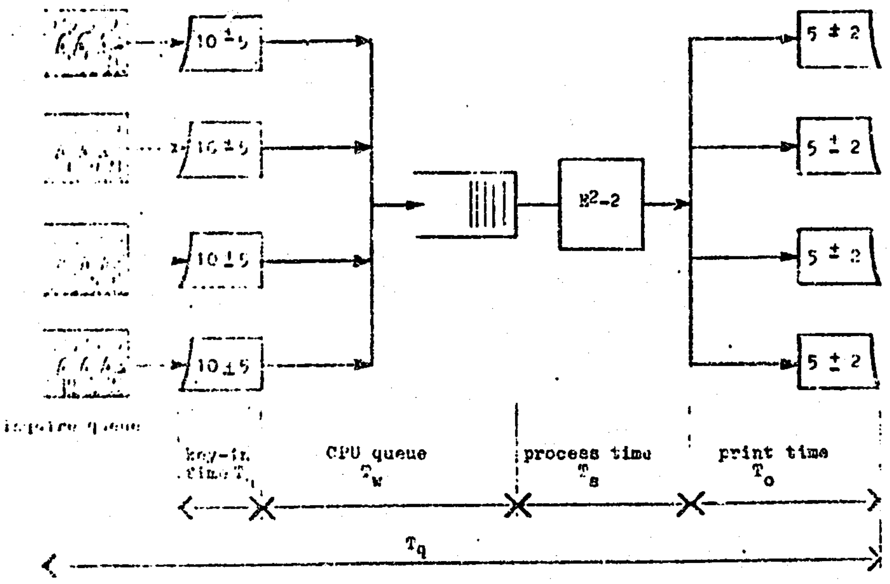

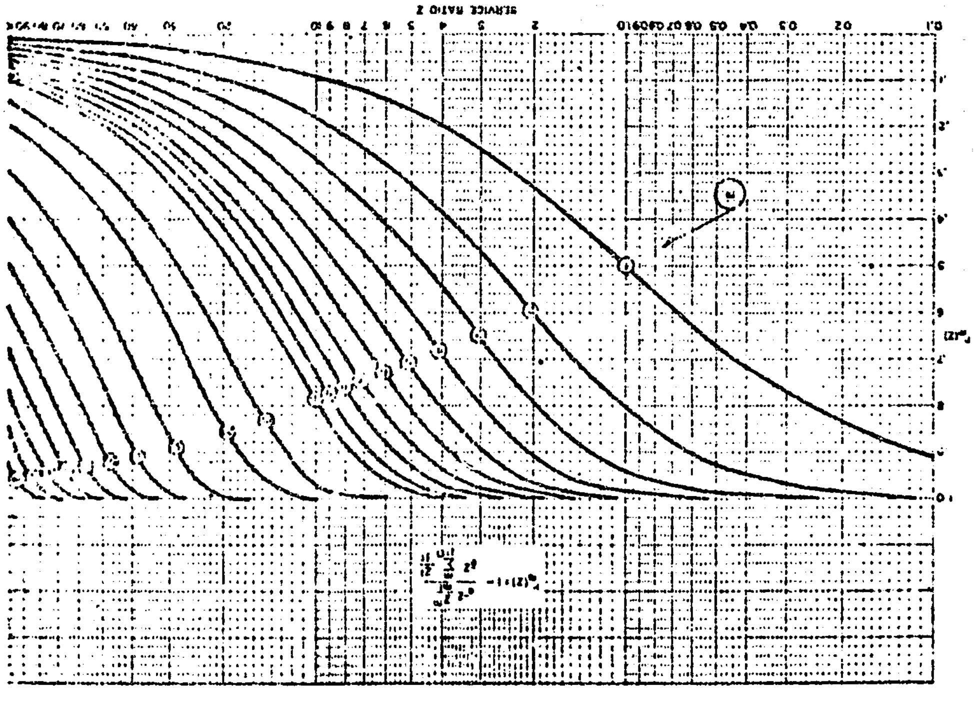

FIGURE 2

Poisson ratio function $E_M(z)$ versus service ratio z, T_s T_r in the classical machine interference problem.

Source: Analysis of Some Queuing Models, (New York, IBM, Ref. F20-0007-1, 1963), p. 40.

Figure 4
Expected response time versus inquiry rate
(simulation results)

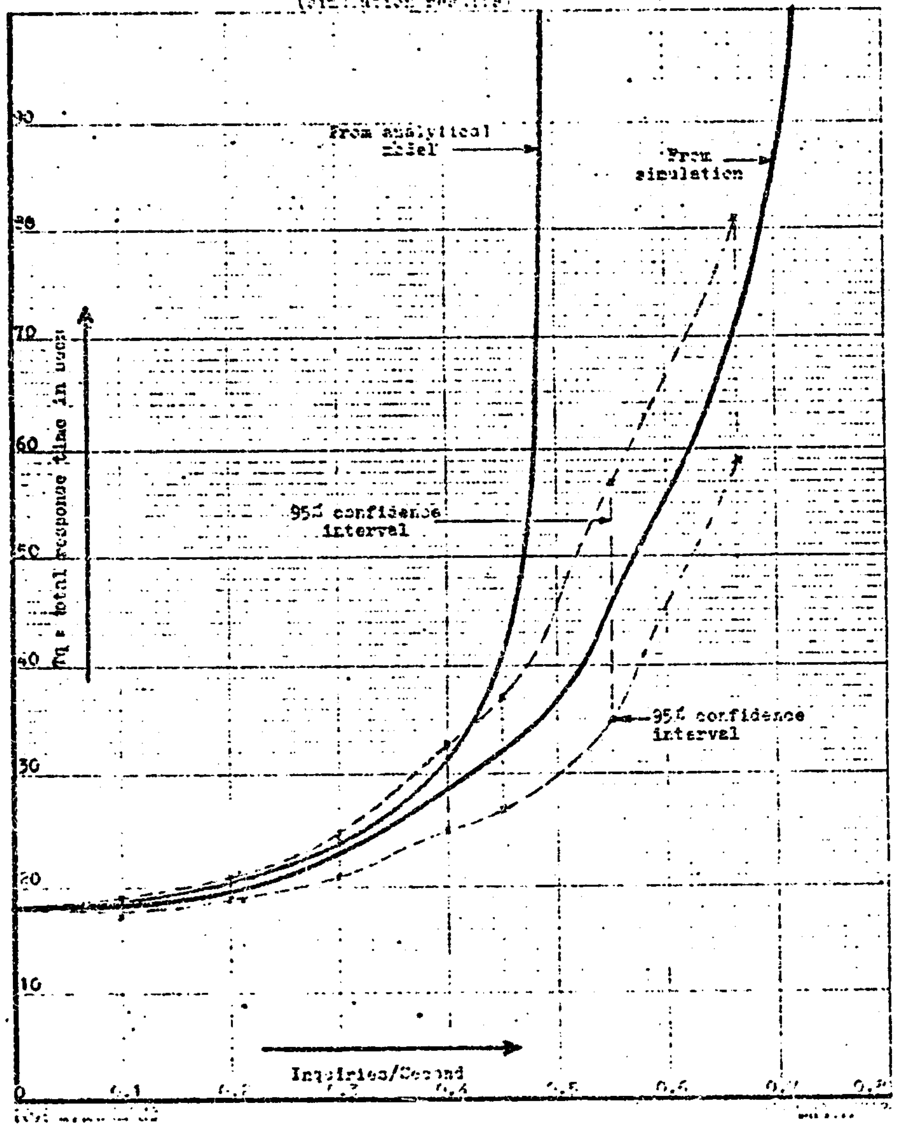

Figure 5
Expected response time versus inquiry rate
(analytical results)