University of Edinburgh

Department of Computer Science

An Experimentally Validated Model of the Paging Drum

by

C. Adams, E. Gelenbe,

J. Vicard

Internal Report

CSR - 10 - 77

September 1977

James Clerk Maxwell Building
The King's Buildings
Mayfield Road
Edinburgh
EH9 3JZ

An Experimentally Validated Model of the Paging Drum

by

C. Adams, E.Gelenbe,

J.Vicard

AN EXPERIMENTALLY VALIDATED MODEL OF THE PAGING DRUM

Colin Adams*,
Department of Computer Science
University of Edinburgh
Edinburgh, Scotland.

Erol Gelenbe,
IRIA-Laboria
78150 Le Chesnay
France.

Jean Vicard,
Département D'Informatique I.U.T.
University Paris Nord
93430 Villetaneuse, France.

ABSTRACT

An experimentally validated paging drum model is presented. The validation is carried out using a simulator driven by a real trace, as well as using an artificial workload.

* Work supported by Science Research Council Grant BRG 93960

1. INTRODUCTION.

In paging systems in a multiprogramming environment, the paging disk (PD) plays an important part. Measurements directly taken on existing systems show that the saturation of the paging channel is often the cause of poor performance. That is why a lot of research has been done on his device's performance.

Coffman [1] has given a mathematical model of the PD assuming that requests for transfer follow a Poisson process. Gelenbe-Lenfant-Potier [2] studied the case in which transfers are sets of grouped and contiguous pages of the PD. In [3] is studied the problem of the simultaneous managements of the main memory of the disk. Baudet-Ferrie-Boulanger [4] considered the case of requests for grouped transfers scattered at random on the PD circumference, assuming the arrival times are given according to a Poisson process.

The Poisson assumption for arrival times has been shown to be unacceptable by measurement experiments: in this paper we report measured squared coefficients of variation (SCV) of interarrival times close to 2 (with the Poisson process, the same parameter is equal to one). The use of diffusion models [8] has been proposed so as to treat the case of non-Poisson arrivals, and some simulation results have been obtained.

The purpose of this paper is to compare results obtained by theoretical models to measured values, to results obtained by simulation general assumptions and to obtain an experimentally validated PD model. The measurement data used is obtained from the EMAS system [14,15] .

1.1. The system being studied

It is a fixed head disk composed of concentric tracks which are divided into equal sectors corresponding to a page of the main memory. The switching of reading to writing can be done in the interval between the passage of the end of a sector and the passage of the beginning of the next sector. A detailed description of this operation can be found in [11] .

The PD can be considered as a queueing system receiving requests defined by a stochastic process describing interarrivals. This process is assumed to be stationary. λ is the average arrival rate. K^2 is the SCV of interarrival times. The system fulfills requests at the maximum rate of one per rotation and per sector (for one track). The load factor of the PD is defined as $\rho = \lambda T$ in which T is the PD rotation time.

In order to increase the PD's performance, a queue is associated to each sector, rather than a single server queue per track. As was done in [1] we will study a different and yet equivalent representation of the physical model. In order to do so we assume that the PD is fixed and that the read/write heads turn around it at constant speed (Figure 1): when each read/write head comes in front of a sector whose queue is not empty, the corresponding transfer request is satisfied.

The sector load factor is defined by $\rho_k = \lambda_k T$ in which λ_k is the average rate of arrivals of requests to the sector k. At statistical equilibrium we must have $0 \leq \rho_k < 1$ and $0 \leq \rho < N$.

1.2. Summary of theoretical results

1.2.1: Poisson arrivals [1].

The mean length of a sector queue is:

$$L_k = \rho_k \left[\frac{N+2}{2N} + \frac{\rho_k}{2(1-\rho_k)} \right] \quad (1)$$

where N is the number of PD sectors.

The steady-state probability that the queue is empty is:

$$P_k(0) = (1-\rho_k) \; e^{-\frac{(N+1)\rho_k}{N}} \left(\frac{1-e^{-\rho_k}}{\rho_k} \right) \quad (2)$$

These results can be obtained through the analysis of the Markov chain imbedded in the process representing the number of customers in the queue, and have been obtained by Coffman [1].

1.2.2. Non-Poisson arrivals [7,8].

Exact analytical results for the paging drum queue with non-Poisson arrivals are unavailable. Therefore it is of interest to investigate the use of an approximate expression for the queue length distribution. We first recall a result obtained in [13]. Independent and identically distributed interarrival times are assumed.

Theorem [13].

Let W_k be the steady-state waiting time of the transfer request arriving to the k-th sector queue. Denote by V the steady-state waiting time of a customer arriving to the GI/D/1 queue (generally distributed independent arrivals and constant service) of service time T and with the same arrival process as to the PD sector. W_k and V are random variables which satisfy the following equality:

$$W_k = V + \eta \quad (3)$$

where η is a random variable uniformly distributed on $[0,T]$ and independent of V. The equality is to be understood in the sense of the equality of the probability distribution functions.

We shall not reproduce here a proof of this result which is given in [13] in a more general context. Its significance lies in the fact that it allows us to compute W_k directly from V if the probability distribution function of V is known, which is not the case in general. We call upon diffusion approximation [6,7] to compute statistics of interest related to V.

Result [7]

The diffusion approximation to the average queue length of a GI/D/1 queue with service time T, arrival rate λ_k and SCV of interarrival time K_k, is given by

$$\hat{L}_k = \rho_k \left(\frac{1}{2} + \frac{\rho_k K_k}{2(1-\rho_k)} \right) \quad (4)$$

Returning to (3), notice that the response time (waiting time plus service time) at the paging drum is $W_k + T/N$, so that the average queue length will be (using (3) and Little's formula):

$$L_k = \lambda_k \; (EV + E\eta + T/N) \quad (5)$$

where E denotes the expectation. But $E\eta = T/2$, therefore

$$L_k = \rho_k \left(\frac{1}{2} + \frac{1}{N} \right) + \lambda_k \, EV \quad (6)$$

But, again using Little's formula and (4)

$$\hat{L}_k = \lambda_k \quad (EV + T)$$

since the service time for the GI/D/1 queue is T. Therefore we have the approximation

$$\lambda_k \, EV \approx \frac{\rho_k^2 K_k}{2(1-\rho_k)} - \frac{1}{2}\rho_k \quad (7)$$

so that we have the approximation we seek:

$$L_k \approx \rho_k \left[\frac{1}{N} + \frac{\rho_k K_k}{2(1-\rho_k)} \right] \quad (8)$$

The expression (8) will be called the diffusion approximation to the average queue length at the PD.

2. EXPERIMENTS: ARTIFICIAL ARRIVAL PROCESS

Experiments have been conducted using large squared coefficients of variation for interarrival times, since it is well established that in such cases diffusion models give less precise results. The simulator itself has been validated using the theoretical results obtained by Coffman [1], resorting to a Poisson process for the arrival sequence. We considered N = 8 sectors and assume the requests are uniformly distributed on these sectors, so the probability for a request to call sector k is $P_k = .125$ for each k. In section 3 we shall discuss some simulation experiments run with arrivals obtained for measurements on an actual operating system.

2.1. The hyperexponential law

The probability density used for interarrival times to the PD is chosen as the sum of two exponential laws, the first taking effect with probability a, the second with probability (1-a). The table below gives for a fixed value of ρ, values of the ratio

λ_1/λ and λ_2/λ as a function of K^2 where λ_1 and λ_2 are the parameters of the two exponential laws, respectively, and $1/\lambda$ is the mean interarrival time.

k^2	a	1 - a	λ_1/λ	λ_2/λ
1	1	0	1	0
2	0.667	0.333	0.677	33.8
6	0.286	0.714	0.296	73.9
16	0.118	0.882	0.128	95.7
50	0.039	0.961	0.049	128.8

Table 1

Increasing K_k^2 implies that we

- increase the block size ((1-a) increases)
- decrease the average time interval between two requests of the same block (λ_2 increases)
- increase the average time interval between two blocks.

This situation is representative of the behaviour of paging systems. For example at the end of a program run an important sequence of transfer requests takes place in a short interval time and then there can be a relatively long period of time before the next set of requests.

The hyperexponential law is not the only probability law which makes it possible to obtain a large SCV. We chose it because on the one hand it accounts for actual facts and on the other hand the event sequence obtained when using this law is easily obtained from a uniform random number generator.

The relation between the global SCV k^2 of interarrivals and sector SCV K_k^2 , under the assumption of independent selection of sector k with probability P_k is:

$$K_k^2 = 1 + P_k(K^2 - 1)$$

2.2.1. The average queue length

On figures 3,4,5,6 we give results for four values of K_k^2. We can see on: - Figure 3 for $K_k^2 = 2$ an excellent precision of the diffusion formula (8) compared with the simulation results. The difference with the formula (1) is already significant for $\rho_k > 0.3$, showing that the Poisson assumption will yield poor results.

- figure 4 for $K_k^2 = 6.1$, significant difference between the simulation results and the diffusion for $\rho_k > 0.8$
- figures 5 and 6 stress the above phenomen even more strongly
- figure 7, the linearity of the function $L_k = f(K_k^2)$ obtained from curves $L_k=f(\rho_k)$ for different values of ρ is apparent.

2.2.2. The empty queue probability

We can see on Figure 8, that queue L_k might be idle for different values of ρ_k. These curves show that $P_k(0)$ tends to $(1-\rho_k)$ when K_k^2 increases. The queue behaviour gets close to GI/D/1.[1] This phenomenon is clearly expressed again in Figure 9, which represents as a function of K_k^2,

- on one hand the number of idle periods
- on the other hand the average length of the queue period

Thus when K_k^2 increases, the block size increases too, and for each request from the same block, except for the first request, the service time is equal to a PD rotation time.

2.3. Empirical formulation

From the simulation results we can propose the following formula to compute the average length of a queue.

$$\hat{L}_k = \frac{\rho_k K_k^2}{2(1-\rho_k)}$$

This formula gives a satisfactory account of the system behaviour for high SCV (See Figures 4, 5, 6).

2.3.1. Simulator structure and measured values

The simulator is a FORTRAN program describing the PD behaviour. The event sequence first registered on a magnetic tape, is composed of pairs (t_i, adr_i).

t_i : the request arrival time to the PD

adr_i: the address of the sector in which the request is to be registered.

$$\ldots < t_{i-1} < t_i < t_{i+1} < \ldots$$

$$adr_i \in \{1, 2, \ldots, N\}$$

The output event sequence includes all the events which took place during the simulation. We call "event" either a transfer request or the end of service for a sector. For each event sequence we have a triplet (adr_i, ℓ_i, t_i).

adr_i : the sector address on which the event took place

ℓ_i: the queue length after the event

t_i: time at which the event took place.

1) GI/D/1 is the queue with general independent interarrivals and constant service times.

The output event sequence is then analysed using FORTRAN programs specially designed to measure the following values:

- the average queue length with confidence interval
- the queue length probability distribution
- the average waiting period
- the average idle queue period

2.3.2. Assumptions

The aim of the simulation is to obtain information about a stochastic process $L = \{L(t)\}$ $t < 0$, from punctual observations of the values of L at times t_i [9,10]. In particular we are interested in the long-term properties of the process. This is only possible under certain assumptions on the stationarity of X. In particular we assume that the process is wide-sense stationary [10] i.e.

i. - the expected value EL(t) is time independent

ii. - the autocovariance

$$COV[L(t), L(t+h)] = \gamma(h) \text{ for any } t \text{ and } h, \text{ and}$$

$$VAR[L(t)] = COV[L(t), L(t+0)] = \gamma(0)$$

What is more, we make the assumption, to be verified in practical examples, that $\gamma(h) \to 0$ if $h \to \infty$, i.e. that measurements more and more distant in time are less and less correlated. With the hypothesis of a wide-sense stationary process we can in a single simulation calculate the first two moments of the random variable $\lim_{t\to\infty} E\,X(t)$ and so determine the confidence intervals of these moments [9,12].

2.3.3. Queue length confidence intervals[9,12]

In the case of a stochastic process $\{L(t)\}$ we estimate $E\,L(t) = \mu$ (under the hypothesis of wide-sense stationarity) by:

$$\bar{L} = \frac{1}{n} \sum_{i=1}^{n} L_i \text{ where } L_i = L(t_i)$$

$\bar{L}$ is an unbiased estimator of μ because :

$$E(\bar{L}) = E\left[\frac{1}{n} \sum_{i=1}^{n} L_i\right] = \frac{1}{n} \sum_{i=1}^{n} E(L_i)$$

$$= E(L_i) = E(L) = \mu$$

With the Bienayme-Tchebycheff inequality we can find a first confidence interval on $\bar{L}$.

$$\text{Prob}\left[\,|\bar{L} - \mu| < \delta\sigma_{\bar{L}}\right] > 1 - \frac{1}{\delta^2} \quad \forall \delta > 0$$

where $\sigma_{\bar{L}} = \sqrt{Var(\bar{L})}$

Thus we can write

$$\bar{L} - \delta\sigma_{\bar{L}} < \mu < \bar{L} + \delta\sigma_{\bar{L}}$$

For a 95% confidence level we have:

$$1 - \frac{1}{\delta^2} = 0.95 \quad \delta = 4.5, \text{ and}$$

$$\bar{L} - 4.5\,\sigma_{\bar{L}} < \mu < \bar{L} + 4.5\,\sigma_{\bar{L}}$$

If we assume that $\bar{L}$ has a normal distribution, the 95% confidence interval becomes:

$$\bar{L} - 1.96\,\sigma_{\bar{L}} < \mu < \bar{L} + 1.96\,\sigma_{\bar{L}}$$

If we assume that the L_i are not correlated, we can write :

$$VAR(\bar{L}) = \frac{1}{n^2} \sum_{i=1}^{n} VAR(L_i) = \frac{1}{n} VAR(L)$$

Thus

$$\sigma_{\bar{L}} = \sqrt{Var(L)/n}$$

We can estimate VAR(L) with the unbiased estimator

$$VAR(L) = \frac{1}{n-1}\left[\sum_{i=1}^{n} L_i - \frac{1}{n} (\sum_{i=1}^{n} L_i)^2\right]$$

We obtain

$$\bar{L} = \frac{1.96\,\sigma_L}{\sqrt{n}} < E(L) < \bar{L} + \frac{1.96\,\sigma_L}{\sqrt{n}}$$

If we increase the number of samples, i.e. if we lengthen the simulation we can in theory, obtain E(L) with a precision which can be fixed beforehand. In order to obtain independent data, we followed Conway's method [9,10], which we briefly summarize. The sequence of samples L_i are divided into n blocks of consecutive samples $L_{11},\ldots,L_{1B},L_{21},\ldots,L_{2B},\ldots,L_{n1},\ldots L_{nB}$ where B is the number of measurements included in any block. We then take

$$L'_i = \frac{1}{B} \Sigma_1^B L_{iB}, \quad 1 \leq i \leq n$$

and treat the $L'_1,\ldots,L'_n$ as n independent samples. The independence assumption will be approximatly satisfied if B is chosen properly and if the correlation between samples decreases with the distance between the samples. The L'_i are then used in the above formulae (instead of the L_i) to obtain the estimate of EL and the confidence interval.

3. SIMULATION EXPERIMENTS USING DATA COLLECTED ON THE EMAS OPERATING SYSTEM

3.1 Data acquisition

The data used in the simulation was collected on EMAS [14] running on the following configuration :

7/8 M-Bytes core store

1 x 2-M Byte (512 pages) drum (revolution time 20 msec).

The data was acquired via a modified version of the EMAS event trace monitor [15] which recorded the times at which all drum requests were issued, and times at which channel programs were started and completion interrupts received. All monitoring was carried out in software however the overhead of this technique is known to be less than 1%. Time stamps on events were accurate to 8 μ sec.

During this run the system was run under a pure demand paging scheme (abandoning the working set replacement prepaging scheme usually employed) and with post paging (i.e. all pages of a process were written out in a group when a process was removed from the multi-programming set). There were around 15 simultaneous users during the monitored session.

The squared coefficient of variation of interarrival times computed from the data is $K^2 = 1.46$ for one set of data and $K^2 = 8.8$ for another. The results are reported in Sections 3.2 and 3.3.

3.2. Simulation: data with $K^2 = 1.46$

The EMAS drum contains four sectors, and in the simulation results reported on Figures 10 A,B,C,D we have chosen N = 4. The cumulative distribution of interarrival times is shown on Figure 11.

The EMAS data was modified in two ways. Firstly we associated an arrival with any given sector with equal probability 0,25 using a random number generator. The assumption of equal distribution of requests on each sector is reasonable in view of the manner in which the EMAS drum is filled [10] . In our simulation run, due to statistical deviations of the number generator the flow of arrivals was not exactly equally distributed among sectors. On Figure 10 we show the simulated average queue lengths at each sector. Notice that in order to vary the arrival rate we have had to vary the average interarrival times to the drum, though K^2 has been kept constant over different values of λ . Therefore this validation is limited the use of a realistic interarrival time distribution with a variable arrival rate, and with the assumption of independent inter-arrival times.

3.3 Trace driven simulation: data with K^2=8.8

In addition to the simulation results discussed in the previous sections, we have conducted a trace driven simulation with another event trace collected on the EMAS system. This data and the results obtained are summarized on Figure 12.

For each value of t (the total real time corresponding to the event trace) the cumulative value of the arrival rate (number of page transfer requests divided t) and of the squared co-efficient of variation of all interarrival times up to time t, were computed. The data corresponding to Sector 2 (shown on Figure 12) was obtained as follows: each page transfer request of the event trace was assigned at random (using a uniform random number generator) to one of the four sectors of the paging drum.

From Figure 12 we observe that the event trace is not stationary (i.e. time independent): the arrival rate changes with time (i.e. ρ_2 decreases slowly), and K_2^2 appears to increase. The cumulative average queue length at sector 2 (measured from the trace driven simulation) also reflects this time dependence. We also show the average queue length $\hat{L}_2$ computed using the empirical formula (given in Section 2.3) is also shown. $\hat{L}_2$ is computed for each value of t by using the corresponding value of ρ_2 and of K_2^2 . These figures correspond to slightly more than 10,000 events per sector.

Obviously the results shown on Figure 12 show that the empirical formula remains valid when compared to the results obtained from the trace driven simulation.

These results show the validity of the empirical formula even when no particular assumptions (stationarity, or independent and identically distributed interarrival times) are made about the arrival stream.

4. CONCLUSION

In this study we have conducted a study of paging drum performance using all three tools of performance evaluation: an analytical model, a simulation model, and data collected on a real system. Non-exponential interarrival times have been assumed, which is a natural departure from previous studies. We have shown that standard models using Poisson arrival assumptions are in general inadequate for predicting average sector queue lengths and we have provided a simple analytical model which has been validated using simulations. The measurement data has been used to extend the simulations to realistic arrival streams.

It is hoped that this study can be extended in the future to include multiple drum configurations as well as the effect of the software overhead in managing these devices.

1. Coffman, E.G.,Jr. "Analysis of a drum input/output queue under scheduled operation in a paged computer system. "J.ACM 16, 73-90 (1969)

2. Gelenbe, E., Lenfant, J., Potier, D . "Response time of a fixed-head disk to transfers of variable length." SIAM J. Computing 4, 461-473 (1975)

3. Gelenbe, E., Lenfant, J., Potier, D. "Analyse d'un algorithme de gestion simultanée mémoire centrale - disque de pagination." Acta Informatica 3, 321-345 (1974).

4. Baudet, G., Ferrie, J., Boulanger, J. "Analysis of a drum with bulk arrivals." International Computing Symposium, 213-224 North-Holland Publishing Company - Amsterdam (1975).

5. Badel, M., Shum, A.V.Y. "Accuracy of an approximate computer system model" in Modelling and Performance Evaluation of Computer Systems, E. Gelenbe, H. Beilner eds, North Holland, 1976.

6. Kobayashi, H. "Applications of the diffusion approximation to queuing networks : Equilibrium queue distributions." J.ACM 21, 316-328 (1974)

7. Gelenbe, E. "On approximate computer system models". J.ACM, 261-269 (1975)

8. Gelenbe, E. "Paging drum statistics with general interarrival times." Unpublished paper..

9. Fishman, G.S. "Concepts and methods in discrete event digital simulation." John Wiley and Sons, New York (1973).

10. Badel, M. "Quelques problèmes lies a la simulation de modèles de systèmes informatiques. "Thèse de Docteur-Ingénieur, Université de Paris VI (1975).

11. Mémoires auxiliares à disques magnétiques DIAD. CII manuel C90097 98.

12. Le Gall, P. "Convergences des simulations et applications aux réseaux téléphoniques. "AFIRO Monographies de Recherche Opérationnelle n° 7 (1968.

13. Gelenbe, E., Yasnogorodsky, R. "A queue with server of walking type (autonomous service)" - to appear.

14. Whitfield H., Wight A.S. "EMAS - The Edinburgh Multi Access System" B.C.S. Computer Journal, Vol.16, n° 4. 1973.

15. Adams J.C., Millard G.E. "Performance Measurements on the Edinburgh Multi Access System" Proceedings ICS-75 at Antibes, France. 1975.

VALIDATED MODEL OF THE PAGING DRUM

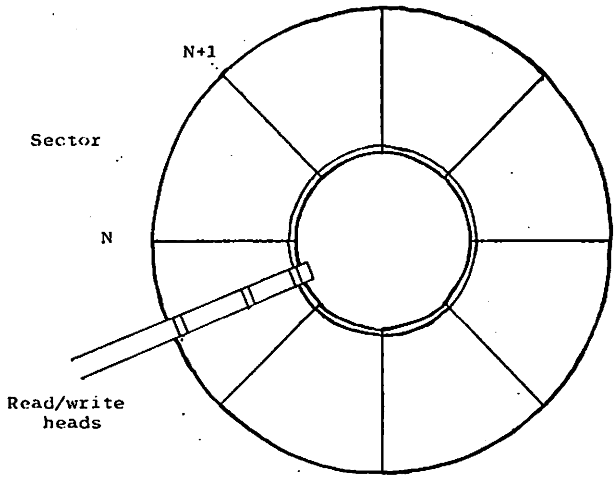

Picture 1 - a The physical model

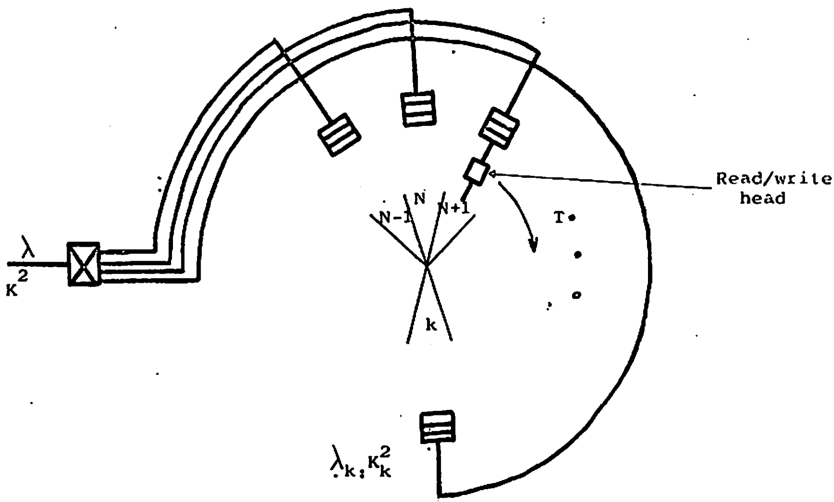

Picture 1 - b. The mathematical model

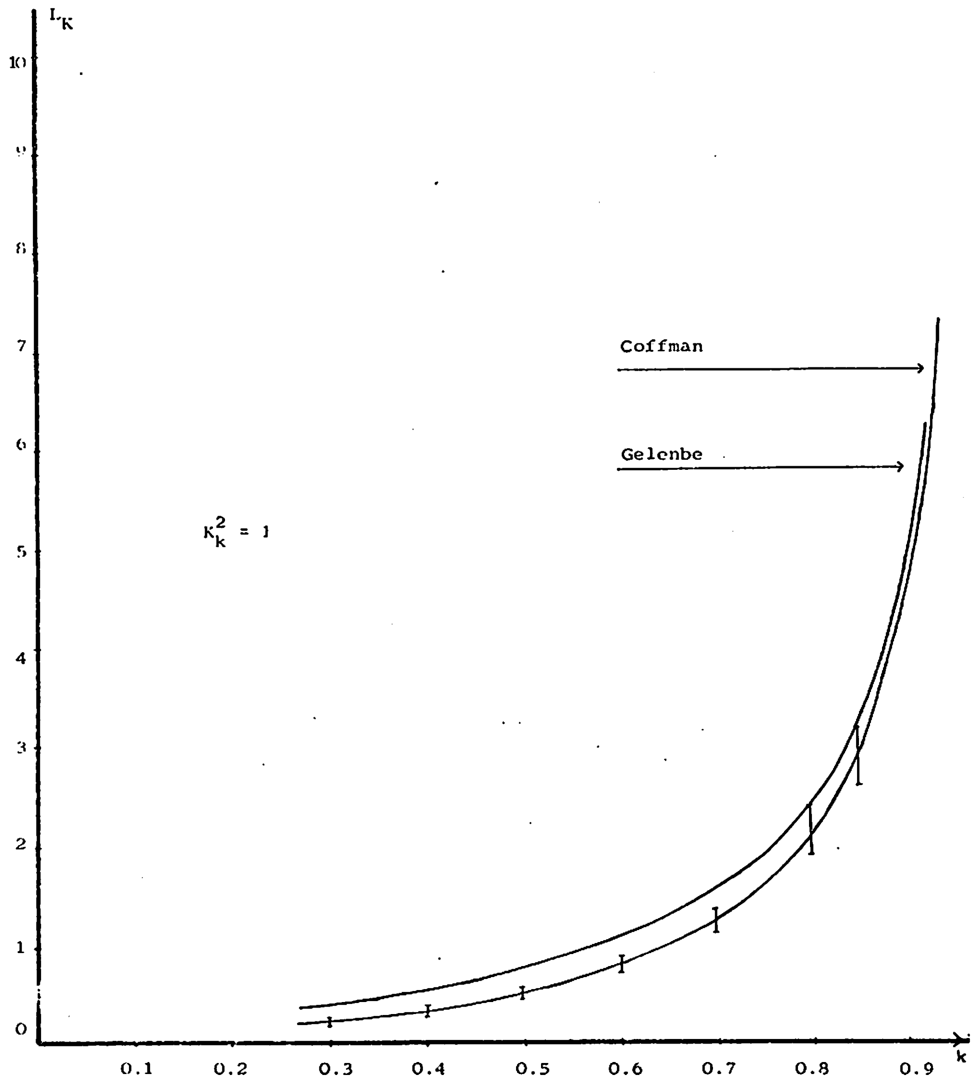

L_K
10
9
8
7
6
5
4
3
2
1
0
Coffman
Gelenbe
$K_k^2 = 1$
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
k

Figure 2.

VALIDATED MODEL OF THE PAGING DRUM

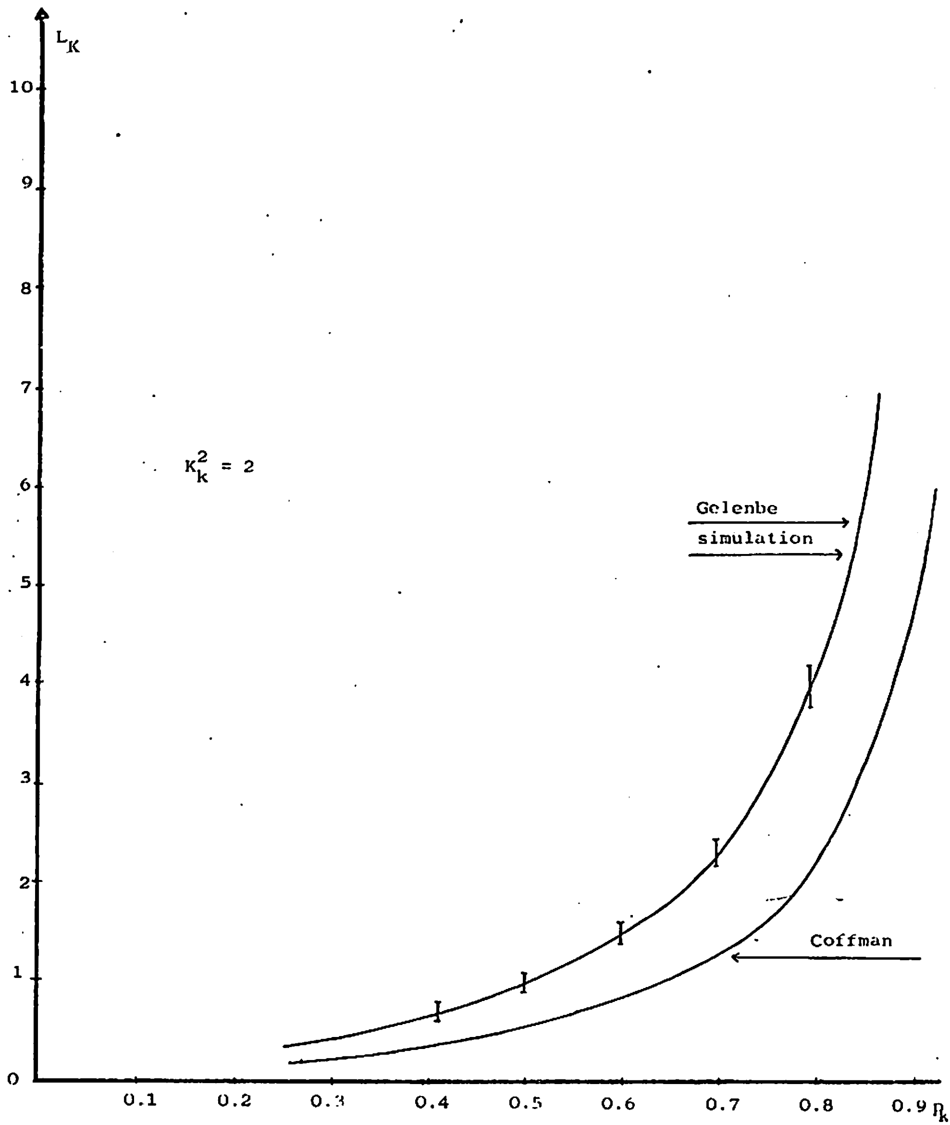

Figure 3.

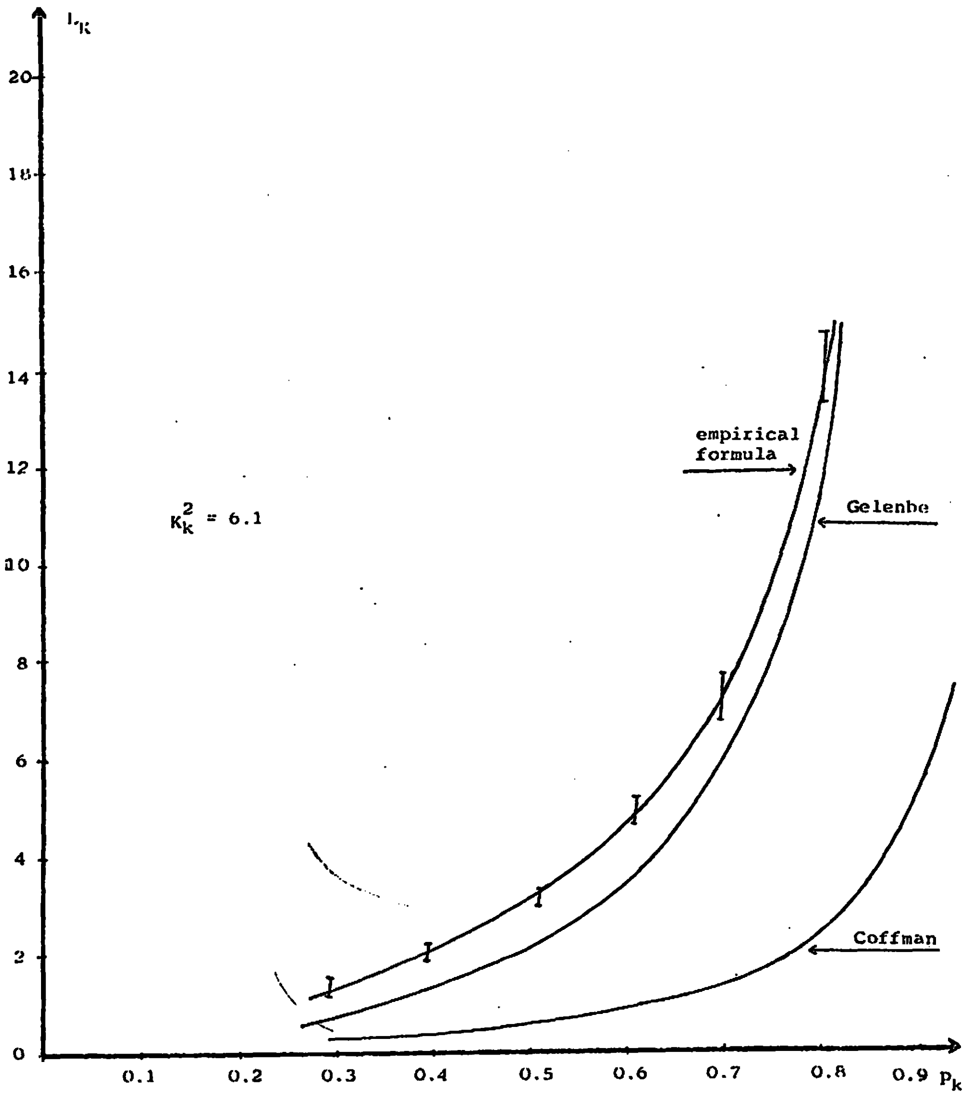
L_K
20
18
16
14
12
10
8
6
4
2
0
$K_k^2 = 6.1$
empirical formula
Gelenbe
Coffman
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
p_k

Figure 4

VALIDATED MODEL OF THE PAGING DRUM

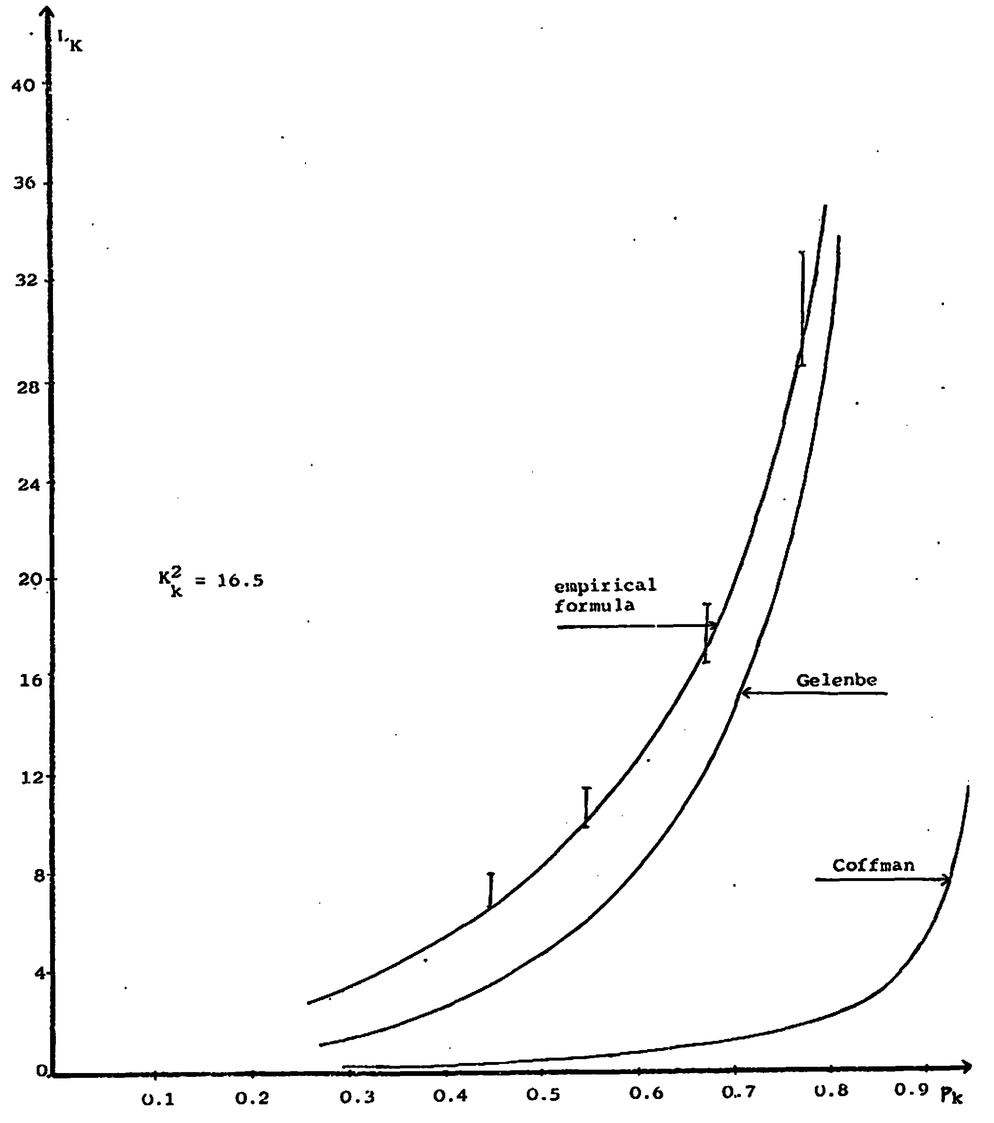

Figure 5.

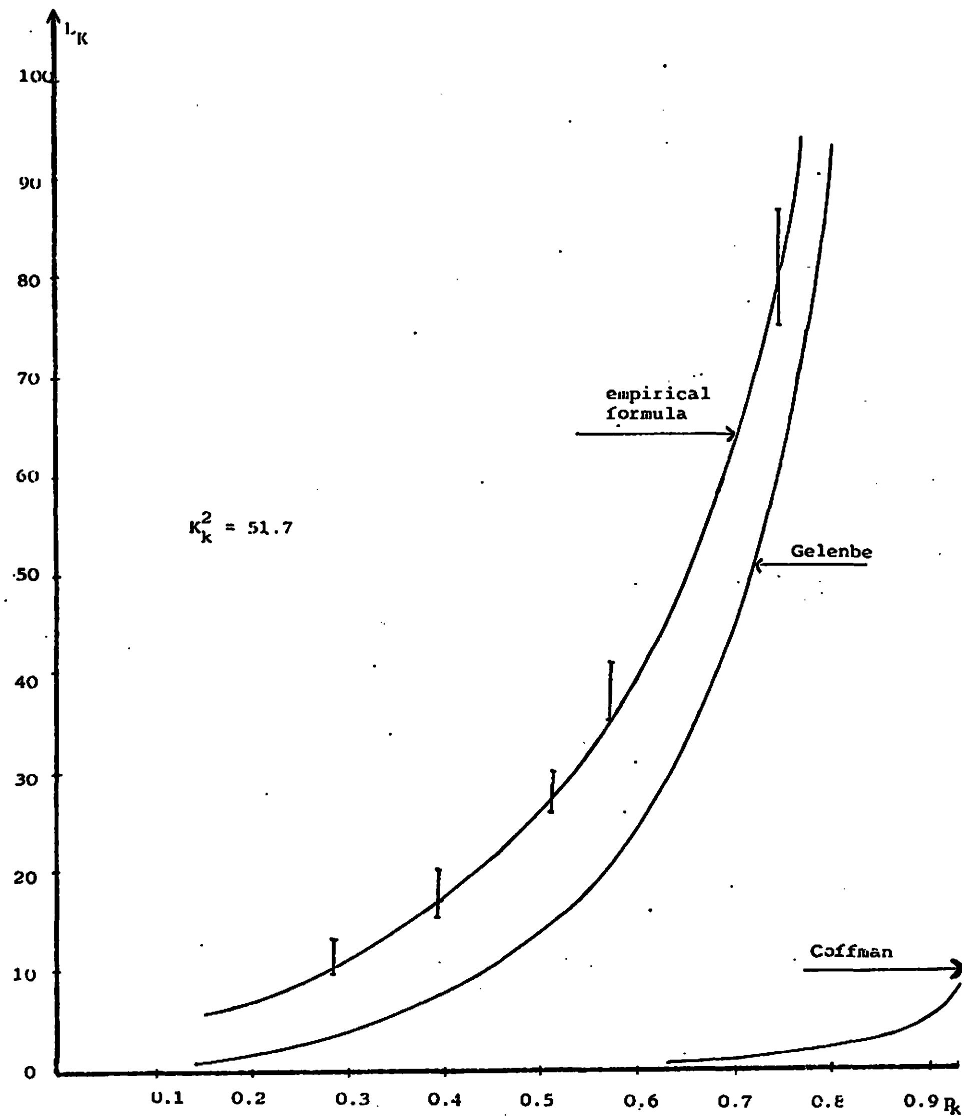
L_K
100
90
80
70
60
50
40
30
20
10
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
ρ_k
empirical formula
$K_k^2 = 51.7$
Gelenbe
Coffman

Figure 6.

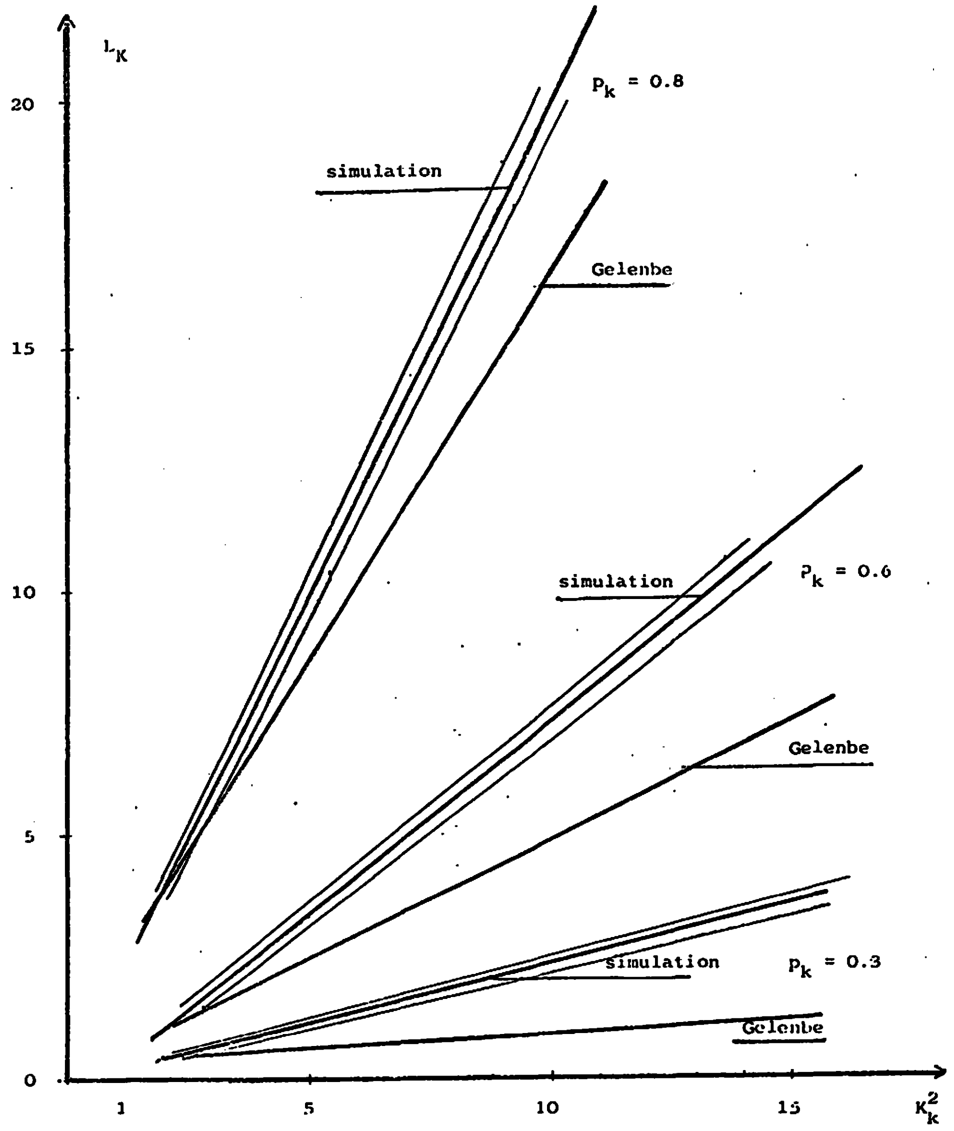
L_K
20
15
10
5
0
1
5
10
15
K_k^2
p_k = 0.8
simulation
Gelenbe
simulation
p_k = 0.6
Gelenbe
simulation
p_k = 0.3
Gelenbe

Figure 7.

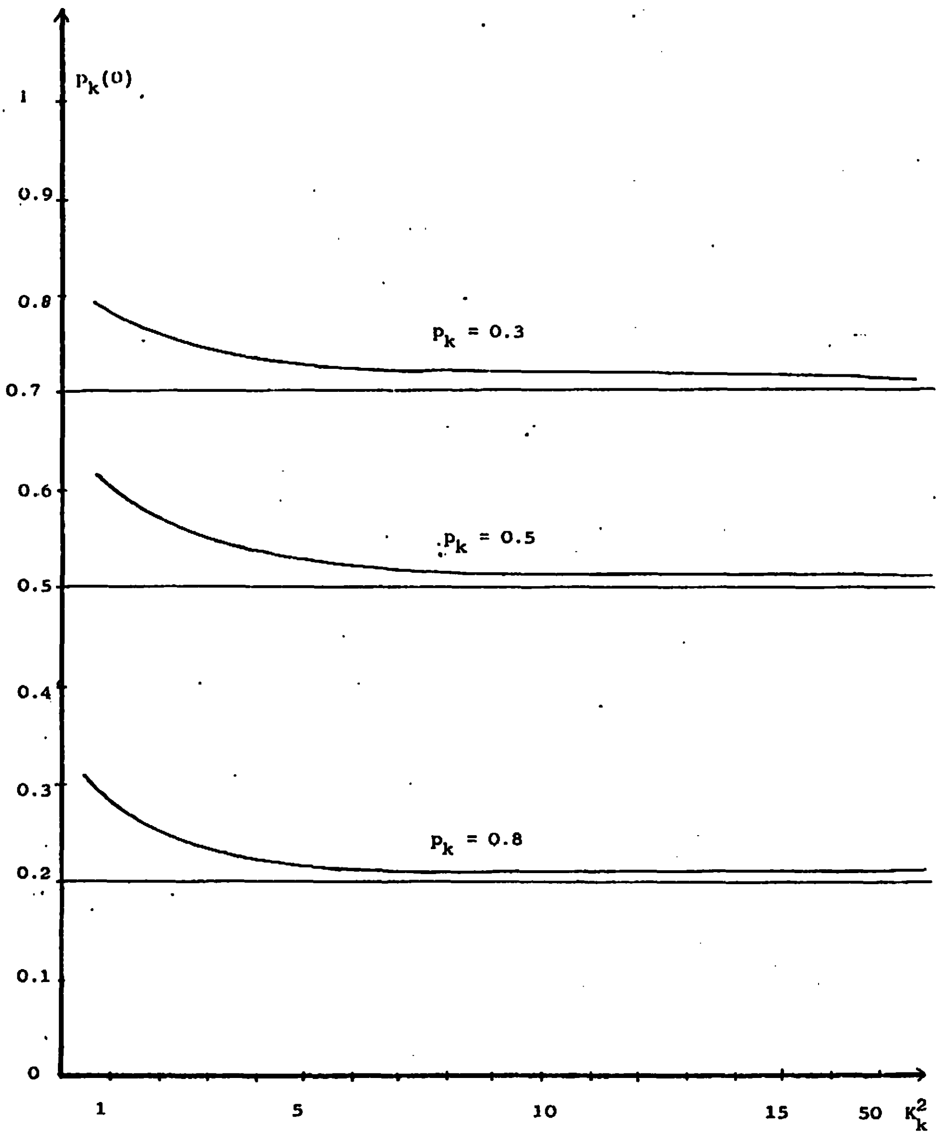

Figure 8.

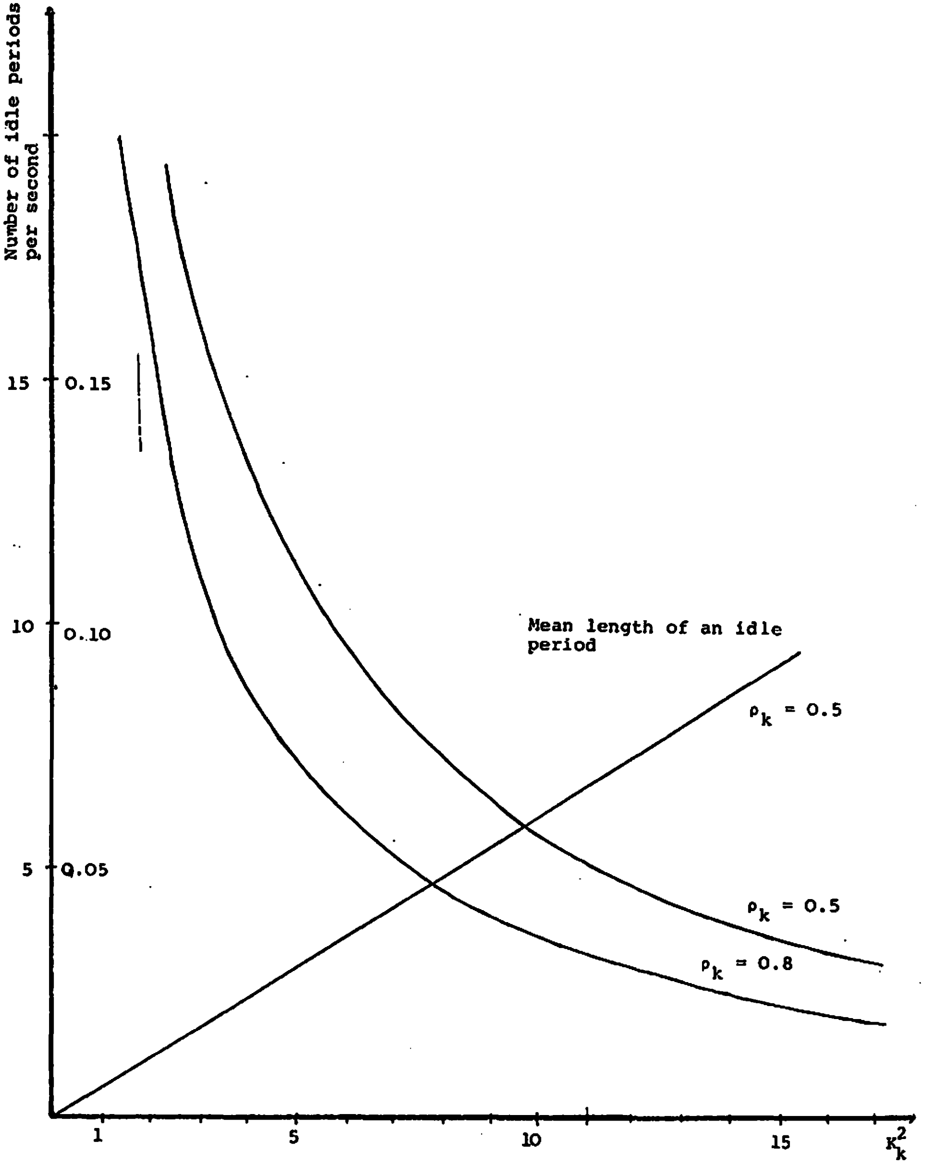
Number of idle periods per second
15
0.15
10
0.10
5
0.05
Mean length of an idle period
ρ_k = 0.5
ρ_k = 0.5
ρ_k = 0.8
1
5
10
15
K_k^2

Figure 9

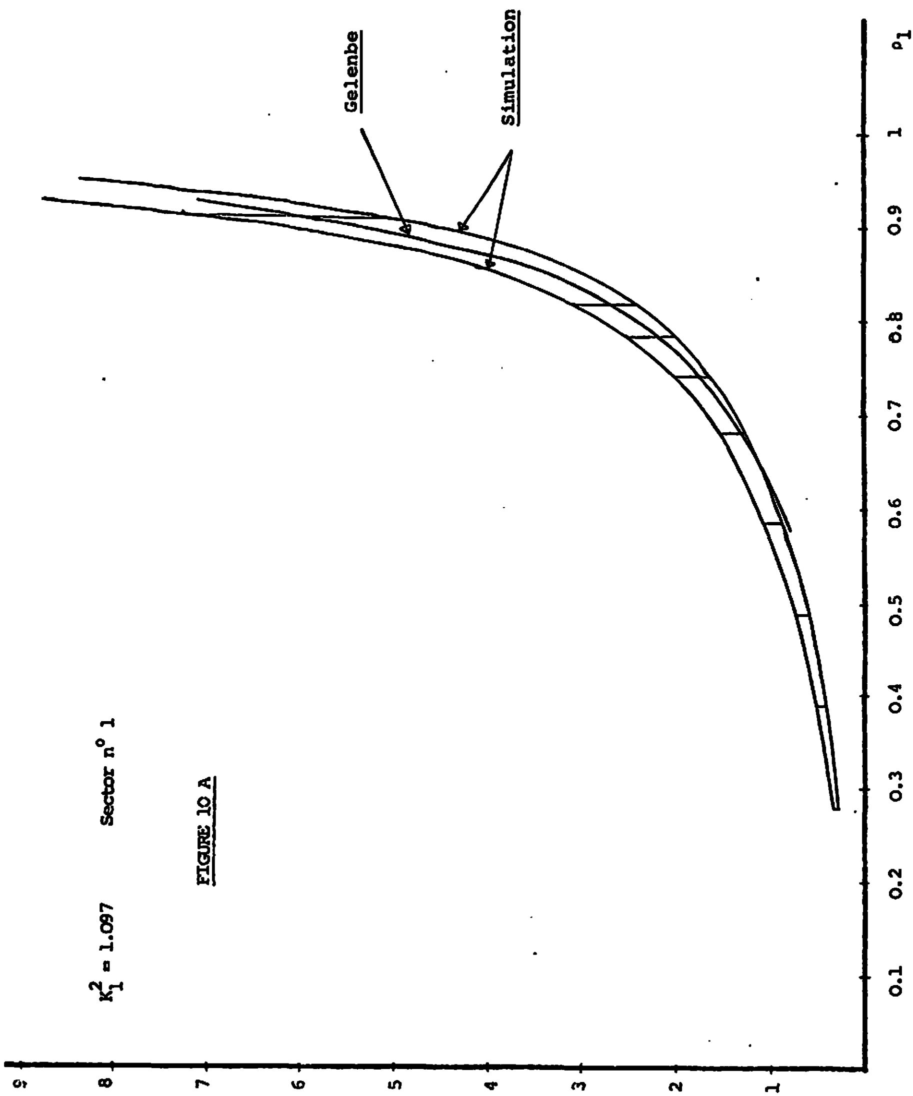
Gelenbe
Simulation
ρ_1
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
1
2
3
4
5
6
7
8
9
$K_1^2 = 1.097$
Sector n° 1

FIGURE 10 A

VALIDATED MODEL OF THE PAGING DRUM

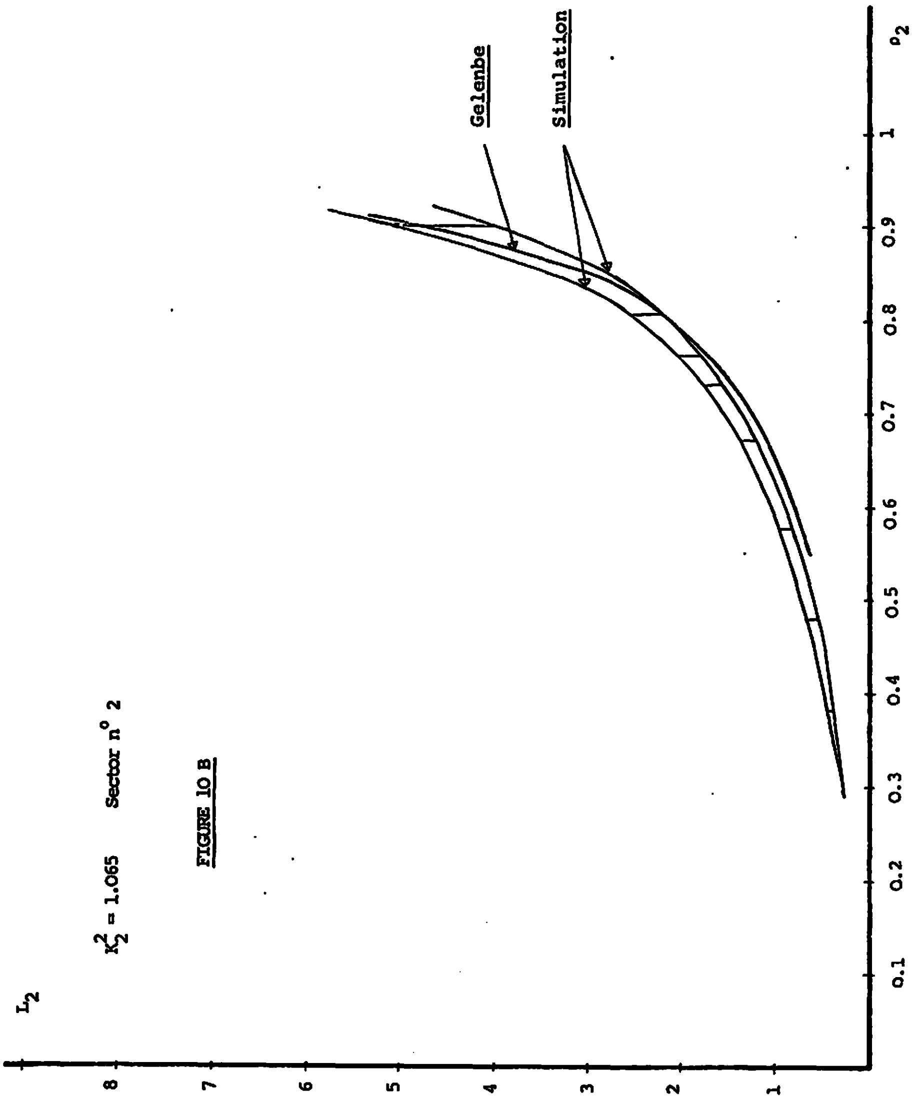

FIGURE 10 B

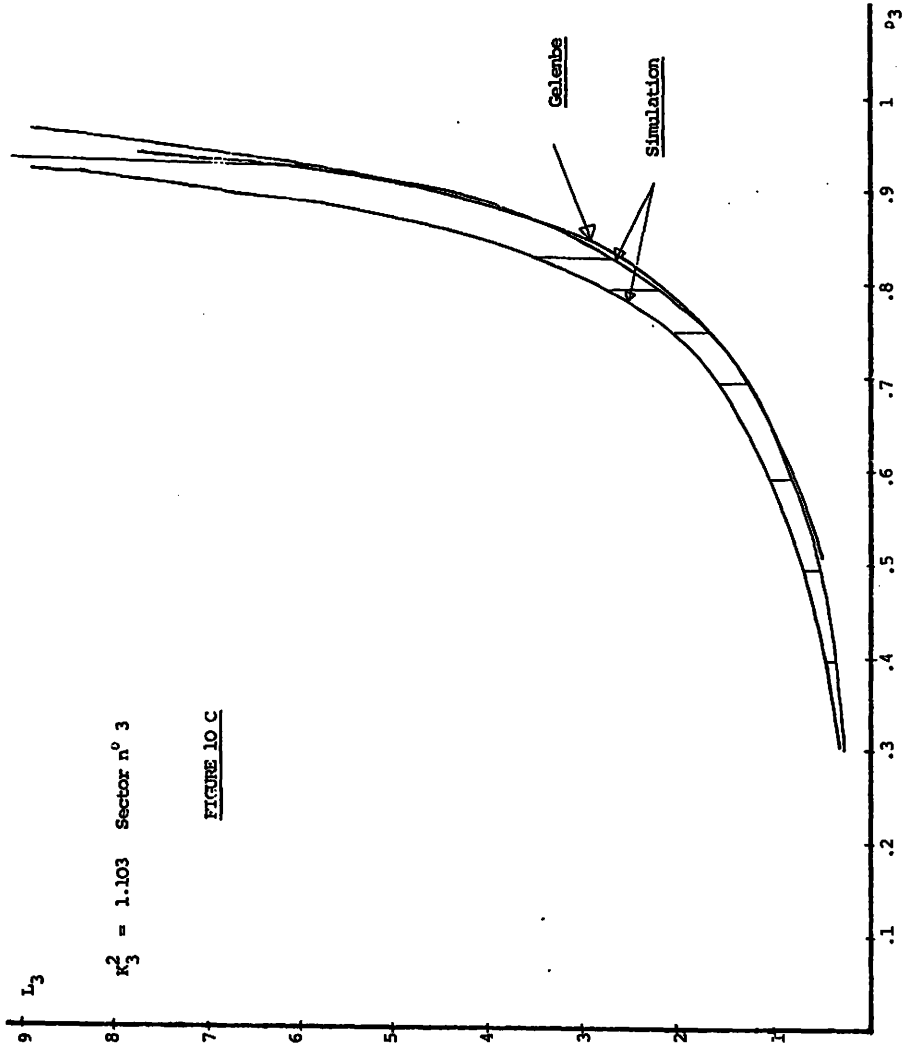
L_3
$K_3^2 = 1.103$ Sector n° 3
Gelenbe
Simulation
9
8
7
6
5
4
3
2
1
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
ρ_3

FIGURE 10 C

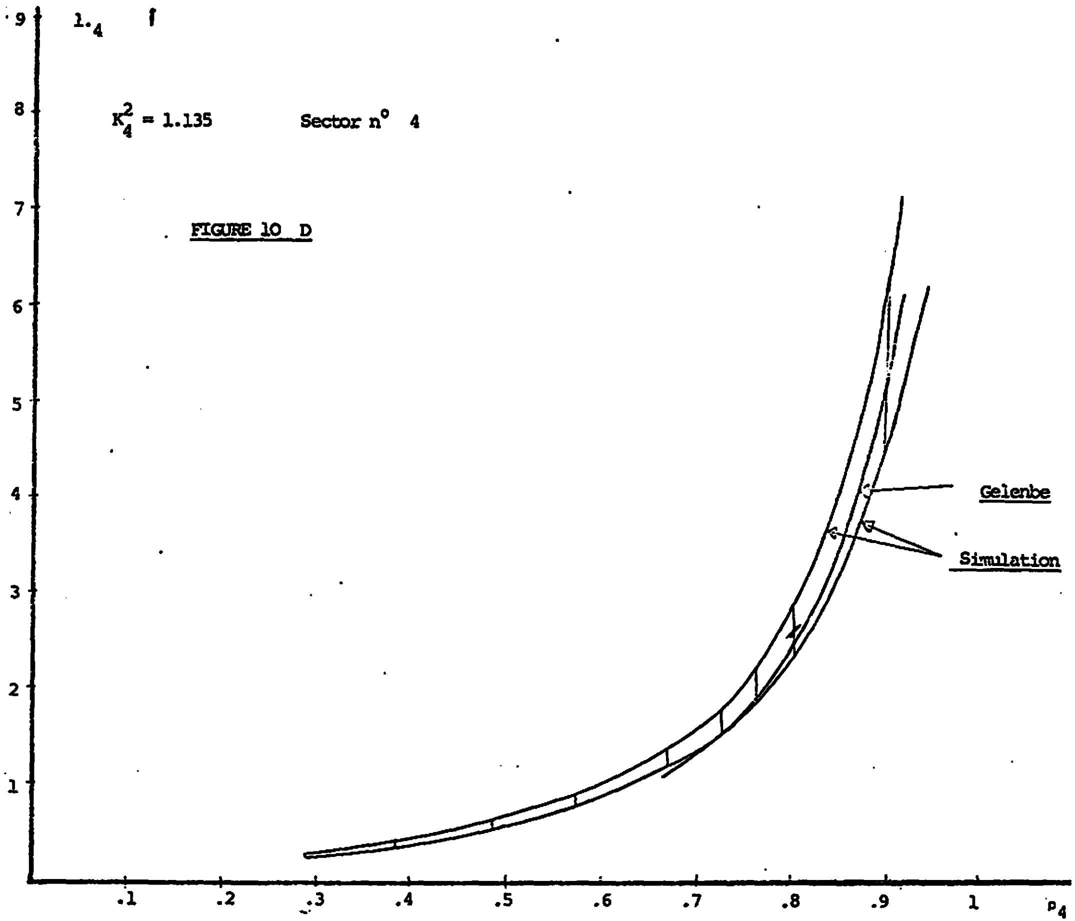
1.4
$K_4^2 = 1.135$
Sector n° 4
Gelenbe
Simulation
9
8
7
6
5
4
3
2
1
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
ρ_4

FIGURE 10 D

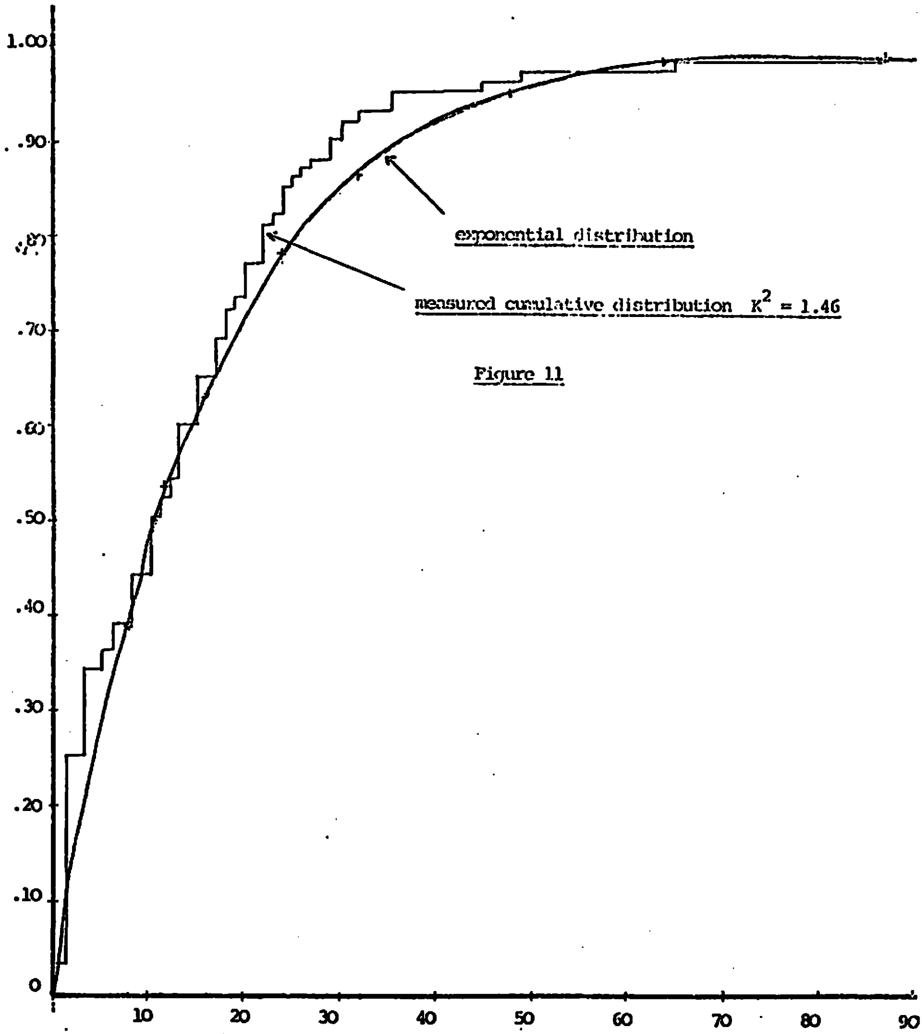
1.00
.90
.80
.70
.60
.50
.40
.30
.20
.10
0
10
20
30
40
50
60
70
80
90
exponential distribution
measured cumulative distribution $K^2 = 1.46$

Figure 11

SECTOR N 2 : TRACE DRIVEN SIMULATION

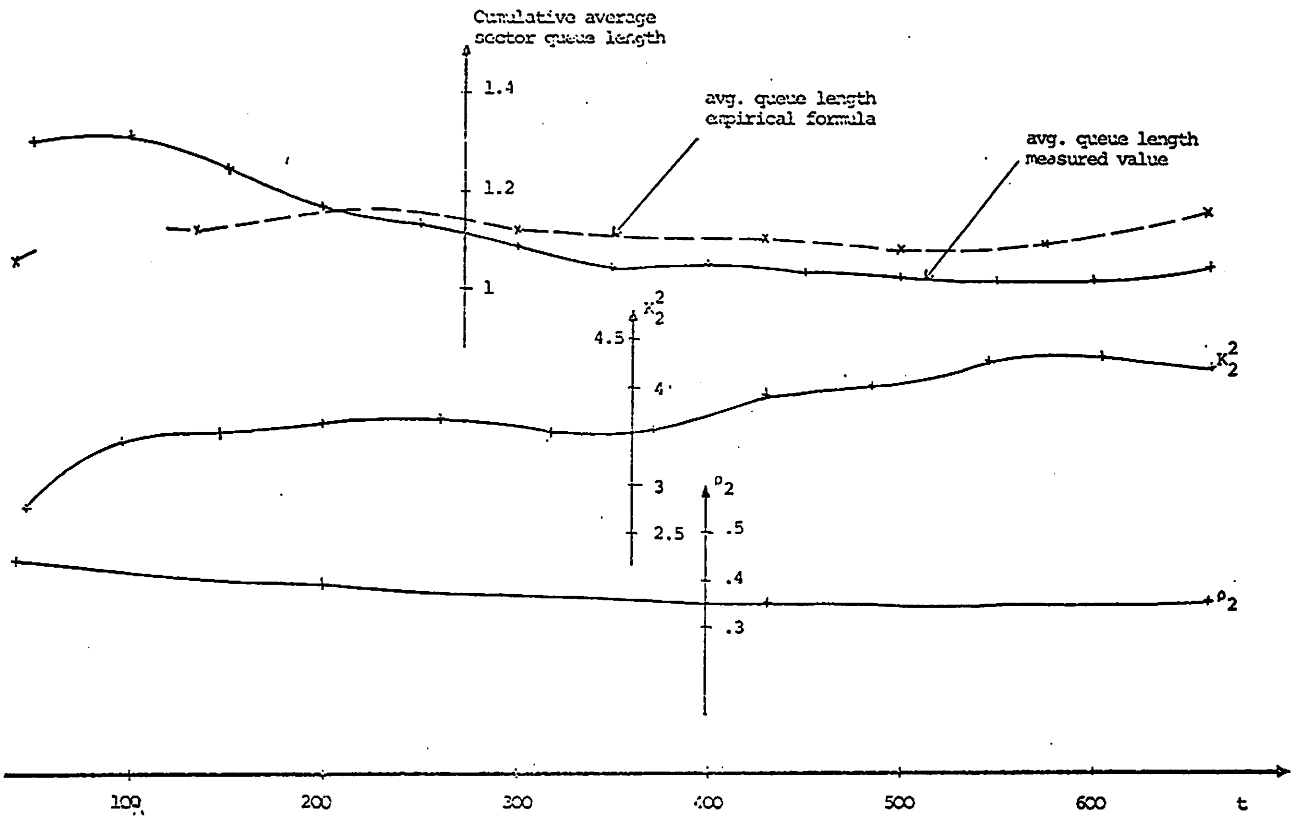

FIGURE 12