SLAC Computation Research Group
Stanford, California

CGTM No. 180
November 1976

PERFORMANCE ANALYSIS OF TWO IO SUBSYSTEMS

Abbas Rafii

Forest Baskett

Computation Research Group
Stanford Linear Accelerator Center

TABLE OF CONTENTS

INTRODUCTION

In this report, we analyze two different IO subsystems using queueing models and simulation techniques.

In the first section, we are concerned with performance analysis of RPS features on disk units like IBM 3330's. Hardware measurements showed high RPS miss ratio during system busy hours. The impact of RPS facility on device performance is investigated and the results are compared with a non-RPS configuration.

In the second section, we consider models of IO subsystems consisting of multiple disk with shared multiple channels. The models are related to the available disk configurations from IBM and ITEL hardwares. The performance of each scheme is analyzed and the results are compared.

This is the second part of a one year project report aimed at performance evaluation of the SLAC Triplex System. The project was supported by helpful discussions with Joe Wells, Dave Folger, Ted Johnston, Richard H. Johnson and Paul Dantzig.

STUDY OF THE PERFORMANCE OF RPS

1. Introduction

The objective of this study is to evaluate the impact of RPS (Rotational Position Sensing) on the response time and utilization of multiple spindle disk drives with a shared channel. Simulation models are used to compare the effectiveness of the RPS scheme with the systems without RPS capability. Analytical models for the number of RPS rotation misses and the utilization of the channel at the saturation point are given.

2. Rotational Position Sensing (RPS) Facility

On some disk drives, like the IBM 3330 and the INTEL 7330's, RPS capability is provided to increase the availability of the control unit or channel by disconnecting each device from the channel during the rotational latency period.

In the RPS scheme, each track is divided [1] into logical sectors. When the Read/Write is positioned on the right track, as the result of an initial seek command, the channel provides the sector number of the data record to be accessed. At this point, unlike the conventional systems, the device is disconnected from the channel and it starts an independent search to find the requested sector. When the sector draws near the R/W head (say, to within two or three sectors before the desired sector), the device sends a signal to the channel. Unless the channel acknowledges this message during the remaining time it takes the sector to reach the R/W head, the sector passes under the head and it will not be available until after one full rotation. The device ready signal is issued each time the sector approaches the R/W arm, and the action is repeated until the channel responds.

3. Measurement Result and Preliminary Analysis

Hardware measurements were done on an IBM 3330 spooling pack at the SLAC Computer Center during different hours of a day. The measurement results show that in peak system usage hours, up to about 50% of the time, the device ready signal is lost due to the busy state of the channel. In such time periods, more than one rotation misses are not uncommon.

A preliminary analysis of the system can bring up some questions. Let 1-C be the RPS hit ratio, i.e., the number of times the channel responds to the device, divided by the number of times the device ready signal is issued. Let T be the time of one full disk rotation. On the average, the effective rotational latency becomes equal to:

$$E(\text{latency}) = (1-C)\,\frac{T}{2} + C(1-C)\left(T+\frac{T}{2}\right) + C^2(1-C)\left(2T+\frac{T}{2}\right) + \ldots.$$

$$= \frac{T}{2} + \frac{TC}{1-C}$$

In the non-RPS system, when the requested track is found and the channel is tied up to a device, the expected time until the beginning of the record is equal to T/2.

In the remaining part of this paper, we compare the actual expected rotational latencies of the two systems.

4. Simulation of RPS Model

Figure 1 is a simulation model of n disk drives which share a comman channel. The model works as follows:

The IO requests arrive at the rate, λ, to a common input queue and are distributed with equal probability among the available devices. The service at each disk consists of a seek stage, followed by a rotational latency stage, and followed by a rotation wait state. When a request completes its rotational latency state, it interrogates the channel status. If the channel is available, the request goes directly to the channel service center. However, if the channel is busy, it has to wait for one full rotation. After one rotation is completed, the channel is interrogated again and the request either goes to another wait state or enters the channel service station.

When a request arrives at the channel, it undergoes service with the service rate equal to μ. The average service time, $1/\mu$, is basically equal to the time to transfer an average size data record. Other assumptions about the model are as follows:

- The devices are identical disk units which work independent of each other.
- A device is busy while its request is in any of the seek, rotational latency, waiting station or channel service station.
- Arrival process is an independent Poisson process with rate λ. Each arrival is a request for an eventual transferring of a record.
- The position of the R/W arm before and after the seek operation is uniformly distributed across the tracks.
- The position of the requested sector (record) is uniformly distributed on the track.
- The average record length is $1/\mu$, measured in units of time to transfer a record. The record sizes have a truncated distribution which limits the size of each record from the above by the size of one track.

6. Results of the RPS Model Simulation

In Tables 1, 2, and 3, the result of RPS model simulation is given for three different mean record sizes. The distribution of record sizes are truncated exponential. One disk rotation time is 18 msec. The maximum seek time is 60 msec.

In Figure 2, the number of RPS rotation misses with respect to the channel utilization is shown. The behavior of the model indicates that as we increase the arrival rate, there is a slow increase on the average queue size. However, after crossing a threshold value, increasing the arrival rate gives rise to a sharp increase in the average queue size. We call this point the saturation point of the system. We also note that beyond this point the number of RPS rotation misses and the channel utilization remain virtually the same. The low channel utilization at the saturation point shows that the system saturates at preparation stages (i.e., seek, rotational latency, wait) much sooner than when the channel is fully utilized.

We note that, as we increase the average record length, the channel utilization at system saturation point increases accordingly.

In Section 8, we compare this result with those obtained in a similar system but without RPS capability.

7. Analytical Models of RPS Performance

Wilhelm [2] gives a detailed analysis of the performance of FCFS (Non-RPS) disks and SLTF (RPS) disks and shows the very small difference in the performance between the two, except at very high arrival rates. In this section, we present simple approximate analytical models for different aspects of our model.

We first find the average number of rotation misses. Let u be the channel utilization. The arrivals at the channel port find the channel busy with probability u. This value is equal to the fraction of arrivals that have to undergo another rotation. Assuming the steady state condition, the latter requests, after completing one rotation, find the channel busy with the same probability and, therefore, u^2 of them has to undergo yet another rotation wait. Similarly, u^3, u^4 ... give the subsequent rotation misses. Therefore, the total number of rotations lost due to the unavailability of the channel is equal to:

$$u + U^2 + u^3 + \dots = \frac{u}{1-u}$$

In the steady state condition, the utilization of the channel is equal to the ratio $\frac{\lambda}{\mu}$, where λ is the arrival rate and $1/\mu$ is the channel service time related to the length of records. We thus have:

$$\text{Number of RPS misses} = \frac{\frac{\lambda}{\mu}}{1-\frac{\lambda}{\mu}} = \frac{\lambda}{\mu-\lambda}$$

This model captures the behavior of the system below the saturation point.

We next find the approximate channel utilization at the saturation point of the system. At the saturation point, we assume all the requests are in the rotational latency as the result of an RPS miss or completion of the seek stage. Let T be time for one rotation. Immediately after the point where the channel has completed transferring a record, it has to wait until the first record reaches the R/W arm on any other drives. If we assume that there are n disk drives, then there are n'<n records waiting to be serviced. If we assume the starting point of the records are uniformly distributed on each track, we can find the average time the channel has to wait until the first record reaches the R/W head.

Let P(t) be the probability distribution function of the time until the first record. Since we assume that the beginning of the records are uniformly distributed in the interval (0,T), we get:

$$P(t) = 1 - \left(\frac{T-t}{T}\right)^{n'}$$

$$E(\text{waiting until first record}) = \int_0^T t\, P'(t)dt$$

$$= \frac{n}{T^n} \int_0^T t(T-t)^{n'-1}dt = \frac{T}{n'+1}$$

Thus, the channel utilization at system saturation is

$$u_s = \frac{1/\mu}{\frac{1}{\mu} + E\ (\text{waiting})} = \frac{1/\mu}{\frac{1}{\mu} + \frac{T}{n'+1}}$$

Since as soon as the transfer of one record is terminated there is a delay until the corresponding device prepares another track due to the seek time period, the value n' in the above result is at least one less than the number of disk drives attached to the channel.

The results obtained from this analytical model are about 5% to 10% off from the results obtained from the simulation (see Table 4). As n grows larger, the result of this analysis becomes more accurate.

8. Simulation of Non-RPS System

In order to compare the effectiveness of RPS on the performance of the disk units, we compare it with a similar system where the RPS feature is not available.

The simulation model of conventional multiple disk organization is shown in Figure 3. In this system, when a seek is completed, the device attempts to obtain the service of the channel. As soon as the channel is available, the device is connected to the channel during the total period of rotational latency and transfer time. Therefore, the busy time of the channel consists of the time until the requested record reaches under R/W arm and the time during which the transfer is executed.

The result of the simulation, under assumptions similar to the first model, is shown in Tables 5 and 6 for two different average record sizes. The channel utilization column is divided into two parts. The first part is the contribution from the rotational delay and the second part is the contribution from transferring of data.

We can see that, compared to the RPS scheme, the system saturates with lower arrival rate. In Figure 4, the average queue length of both models, for different arrival rates, are shown. We can see that for a given arrival rate, the RPS scheme gives lower queue size. From Little's result, the queue length is proportional to the average waiting time of the requests. This shows that RPS system gives better response time than the conventional systems.

These results are also reflected in the channel utilization values due to transfer operation. For a given arrival rate, we get a lower utilization in the non-RPS system. We also note that the non-RPS system saturates with lower arrival rate than the first model. For instance, the ratio of channel utilization due to the transfer operation for RPS over the conventional system is 0.25/0.13 = 1.92 for average record sizes equal to 1.5 msec. The same ratio is equal to 1.6 for record sizes of 4.5 msec.

9. Effect of Other Parameters on the Efficiency of the RPS Scheme

The high variance of record sizes might have adverse effect on the efficiency of the RPS system. Since the record sizes are bounded from the above by the length of one track, for a given average record size we cannot get an arbitrary large variance. We tried a truncated hyper-exponential record length distribution with coefficient of variation equal to about 3. The result didn't show significant degradation of RPS performance.

We can investigate the effect of the number of disk drives attached to the same channel on the efficiency of RPS. The analytical models obtained in Section 7 give a partial answer to this question. Namely, under the saturation point, the number of RPS misses should remain the same and the channel utilization at the system saturation point should increase as the number of disk drives increases. The simulation of the RPS model with six and eight disks, basically gives the same results (Table 7).

10. Conclusion

We have studied the effectiveness of Rotational Position Sensing on the performance of multiple spindle disk drives.

We have seen that the RPS feature does indeed improve the performance of disk storage systems, increases the availability of channel, and gives better effective channel utilization. For a given arrival rate, the RPS scheme gives lower response time than the similar system without RPS facility. When we use the RPS system, the system saturates with higher arrival rate than the conventional schemes.

When we increase the number of disk drives connected to the channel, the number of RPS misses stays virtually the same under the saturation point. However, as we increase the number of disks, the channel utilization at system saturation point increases.

We have also seen that in our model, the RPS system saturates much sooner than when the channel is allowed to be fully utilized.

Another problem which needs further investigation is the effect of access window length on the efficiency of RPS. The access window is the time period in which the channel can respond to a device ready signal before the sector is gone passed the arm. It seems that the RPS efficiency would degrade as we increase access window size.

REFERENCES:

1. "Reference Manual for IBM 3830 Storage Control Model 2", GA26-1617-4.

2. Wilhelm, N.C., "Analysis of Multiple Spindle Disk Drives'" Ph.D. Dissertation, Electrical Engineering Dept., August 1973.

Arrival Per Msec	λ/μ	Av. Q Length	Number of rotations lost due to RPS	Channel Utilization	Number of Times Q becomes Empty
0.05	0.075	0.001	0.085	0.07	1100
0.07	0.105	0.016	0.123	0.11	1100
0.085	0.127	0.051	0.151	0.13	1100
0.10	0.15	0.153	0.178	0.15	1100
0.115	0.172	0.350	0.202	0.17	1100
0.13	0.19	0.84	0.247	0.20	1100
0.17	0.255	754.7	0.300	0.25	748
0.20	0.30	8155.2	0.298	0.25	134

TABLE 1. RPS Simulation Results with Poisson Arrivals
Record Length = $1/\mu$= 1.5 msec

Arrival Per Msec	λ/μ	Av. Q Length	Number of rotations lost due to RPS	Channel Utilization	Number of Times Q becomes empty
0.05	0.15	0.0047	0.189	0.15	1100
0.07	0.21	0.0408	0.281	0.21	1100
0.085	0.255	0.19	0.389	0.26	1100
0.10	0.30	0.47	0.434	0.30	1100
0.115	0.345	1.59	0.522	0.34	1100
0.13	0.39	12.42	0.611	0.39	1100
0.17	0.51	8101.3	0.645	0.41	140
0.20	0.60	15541.0	0.656	0.41	25

TABLE 2. RPS Simulation Results with Poisson Arrivals
Record Length = $1/\mu$ = 3.0 msec

Arrival Per Msec	λ/μ	Av. Q Length	Number of rotations lost due to RPS	Channel Utilization	Number of Times Q becomes empty
0.05	0.225	0.01	0.307	0.22	1100
0.07	0.315	0.10	0.505	0.31	1100
0.085	0.382	0.48	0.635	0.36	1100
0.10	0.450	1.81	0.810	0.44	1100
0.115	0.517	205.0	0.982	0.51	1100
0.13	0.585	3903.0	0.986	0.51	30
0.17	0.765	13365.0	0.988	0.51	12

TABLE 3. RPS Simulation Results with Poisson Arrivals
Record Length = $1/\mu$ = 4.5 msec

	n' = 4	n = 6
Average Record Length	Utilization from Analytical Model	Utilization from Simulation Model
1.5	0.29	0.25
3.0	0.45	0.41
4.5	0.55	0.51

TABLE 4. Comparison of the channel utilization values obtained from the approximate analytical model with simulation results

Arrival Per Msec	λ/μ	Av. Q Length	Channel Utilization Rotational Latency	Transfer Time	Number of Times Q becomes empty
0.05	0.225	0.22	0.43	0.21	1101
0.07	0.315	320.0	0.62	0.30	52
0.085	0.382	3481.0	0.64	0.31	34
0.10	0.45	6829.0	0.65	0.32	15

TABLE 6. Non-RPS Simulation Results with Poisson Arrivals
Record Size = $1/\mu$ = 4.5 msec

Arrival Per Msec	λ/μ	Av. Q Length	Channel Utilization Rotational Latency	Transfer Time	Number of Times Q becomes empty
0.05	0.075	0.04	0.47	0.08	1101
0.07	0.105	0.60	0.60	0.10	1101
0.085	0.127	85.0	0.77	0.128	798
0.10	0.150	2746.0	0.80	0.132	17
0.115	0.172	6288.0	0.81	0.137	17
0.13	0.195	9526.0	0.82	0.137	17

TABLE 5. Non-RPS Simulation Results with Poisson Arrivals
Record Length = $1/\mu$ = 1.5 msec

8 Disk Drives					6 Disk Drives
Arrival Per Msec	λ/μ	Av. Q Length	Number of Rotations lost due to RPS	Channel Utilization	Channel Utilization
0.05	0.15	0.0	0.164	0.14	0.15
0.07	0.21	0.007	0.348	0.22	0.21
0.085	0.255	0.010	0.371	0.25	0.26
0.10	0.30	0.066	0.459	0.30	0.30
0.115	0.34	0.340	0.597	0.33	0.34
0.13	0.39	0.555	0.665	0.38	0.39
0.17	0.51	2306.0	0.971	0.48	0.41

TABLE 7. Comparison of RPS Simulation Results Using 8 and 6 Disk Drives

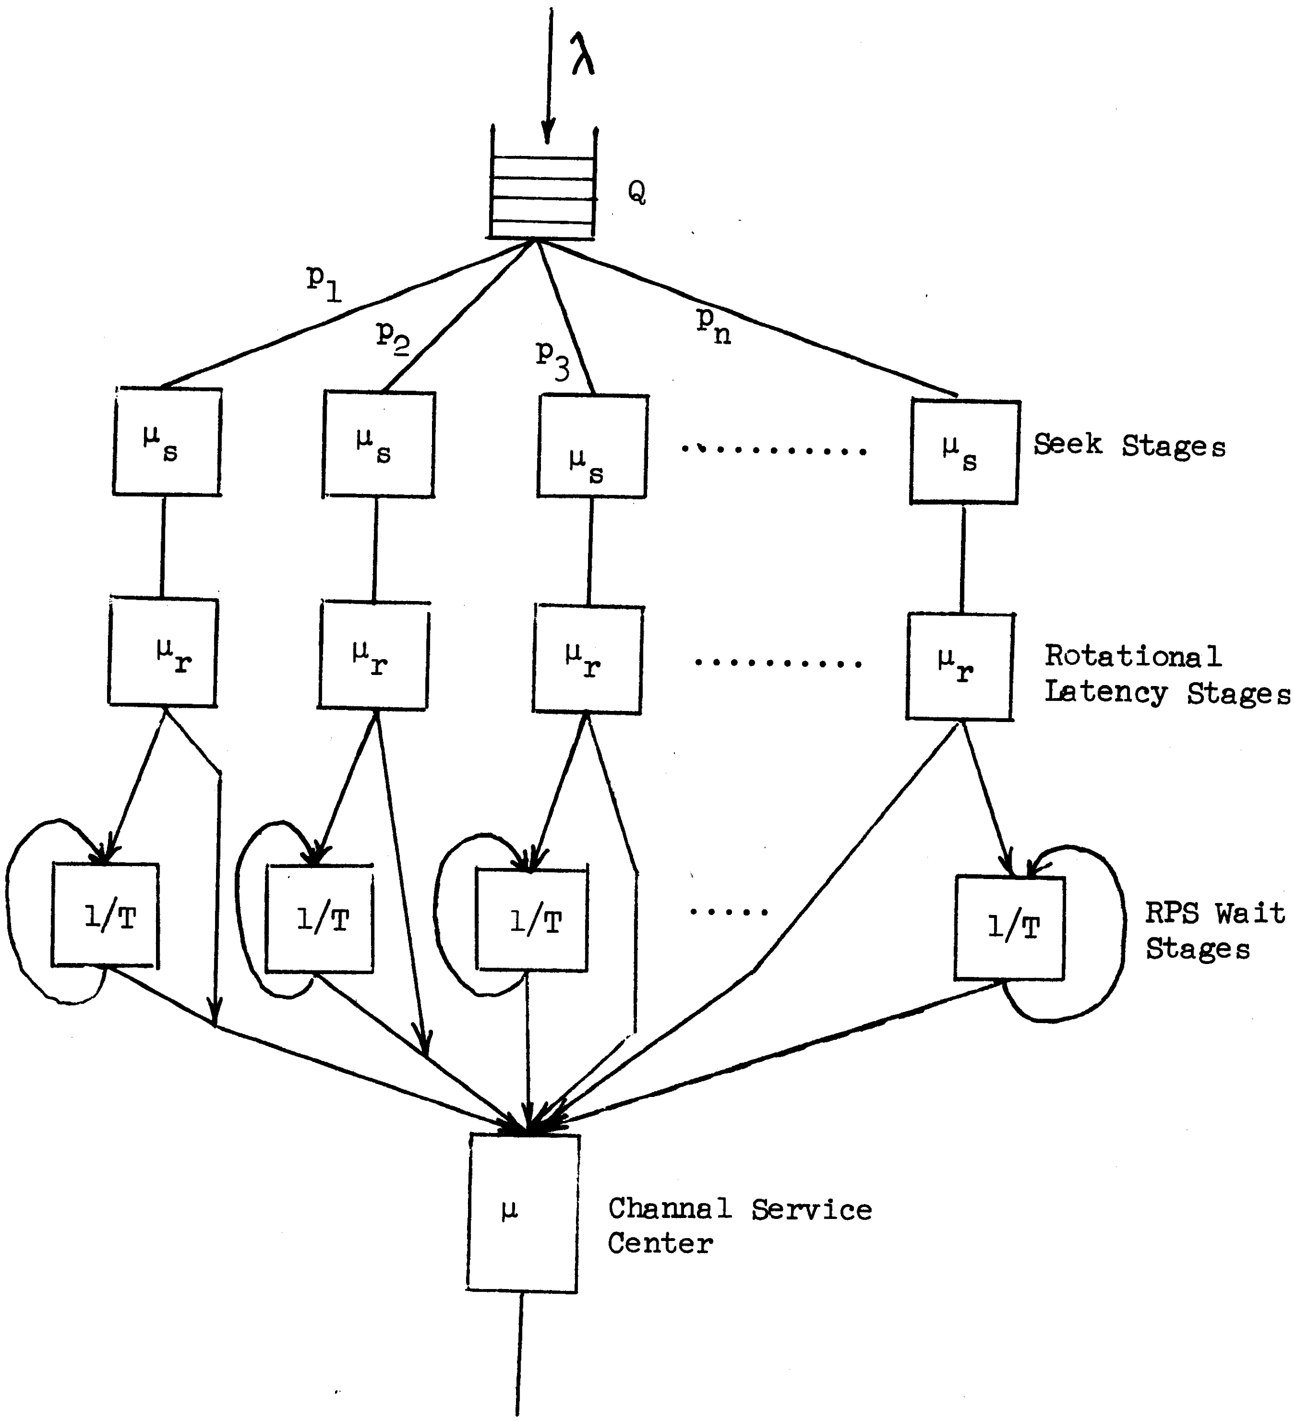

<u>Figure 1:</u> Simulation model of RPS scheme

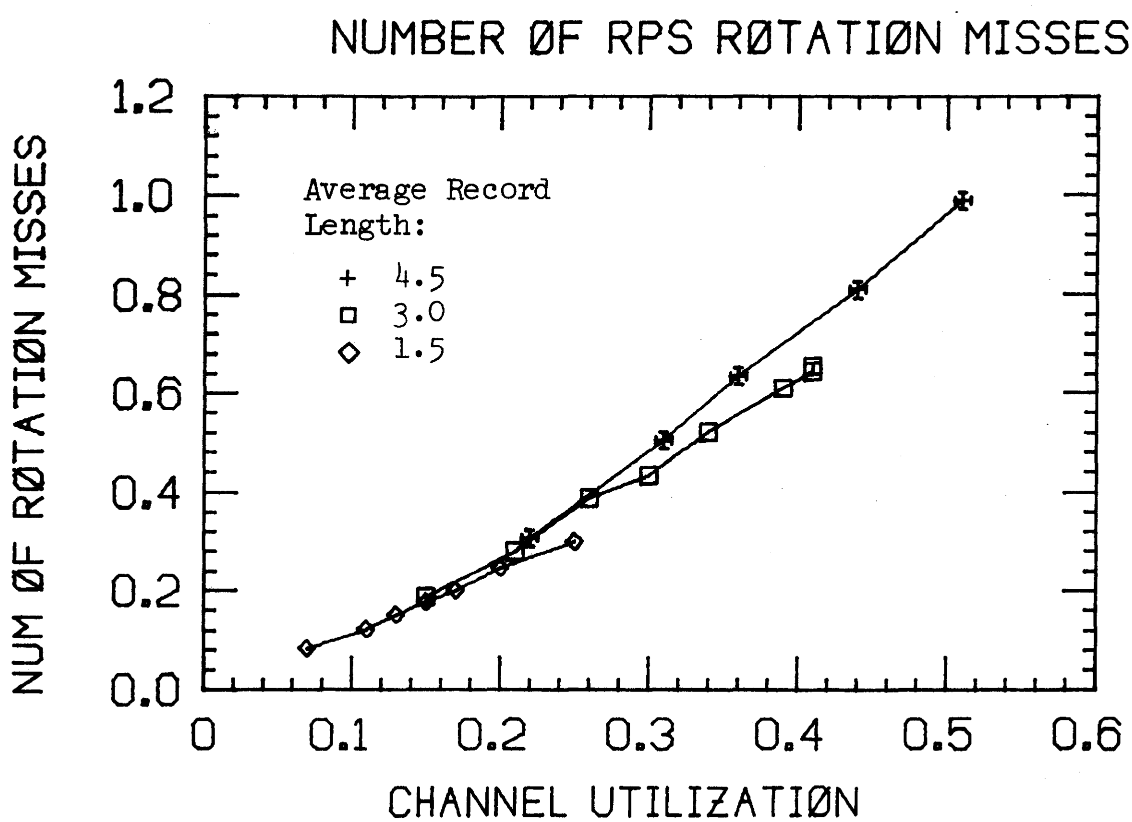
NUMBER ØF RPS RØTATIØN MISSES
NUM ØF RØTATIØN MISSES
CHANNEL UTILIZATIØN
Average Record Length:
+ 4.5
□ 3.0
◇ 1.5
0.0
0.2
0.4
0.6
0.8
1.0
1.2
0
0.1
0.2
0.3
0.4
0.5
0.6

Figure 2

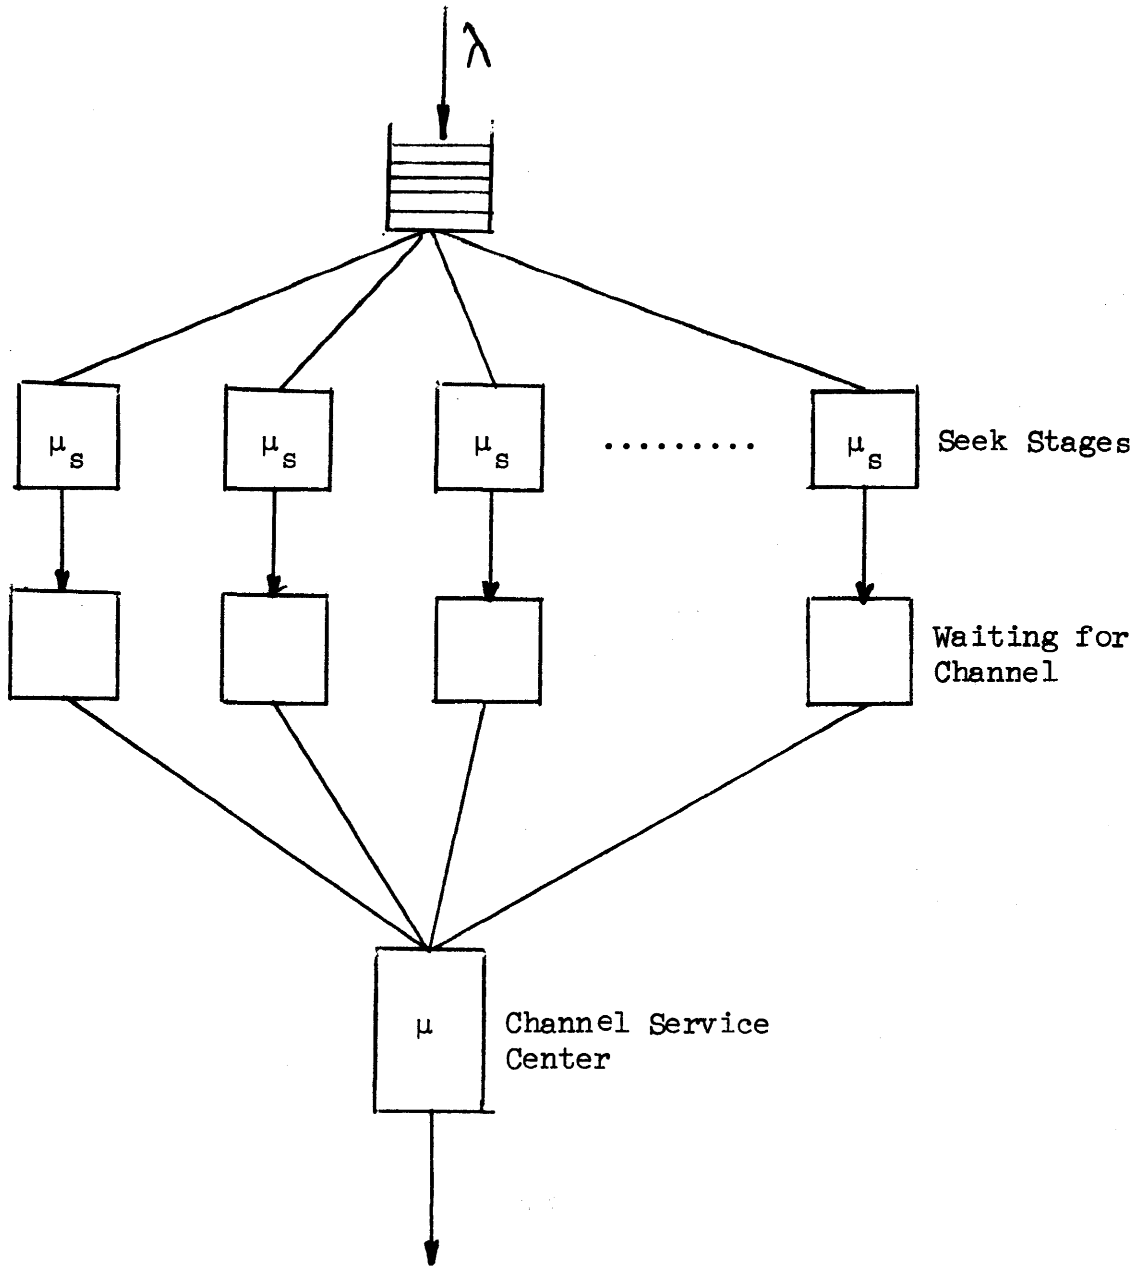

Figure 3: Non-RPS scheme simulation model

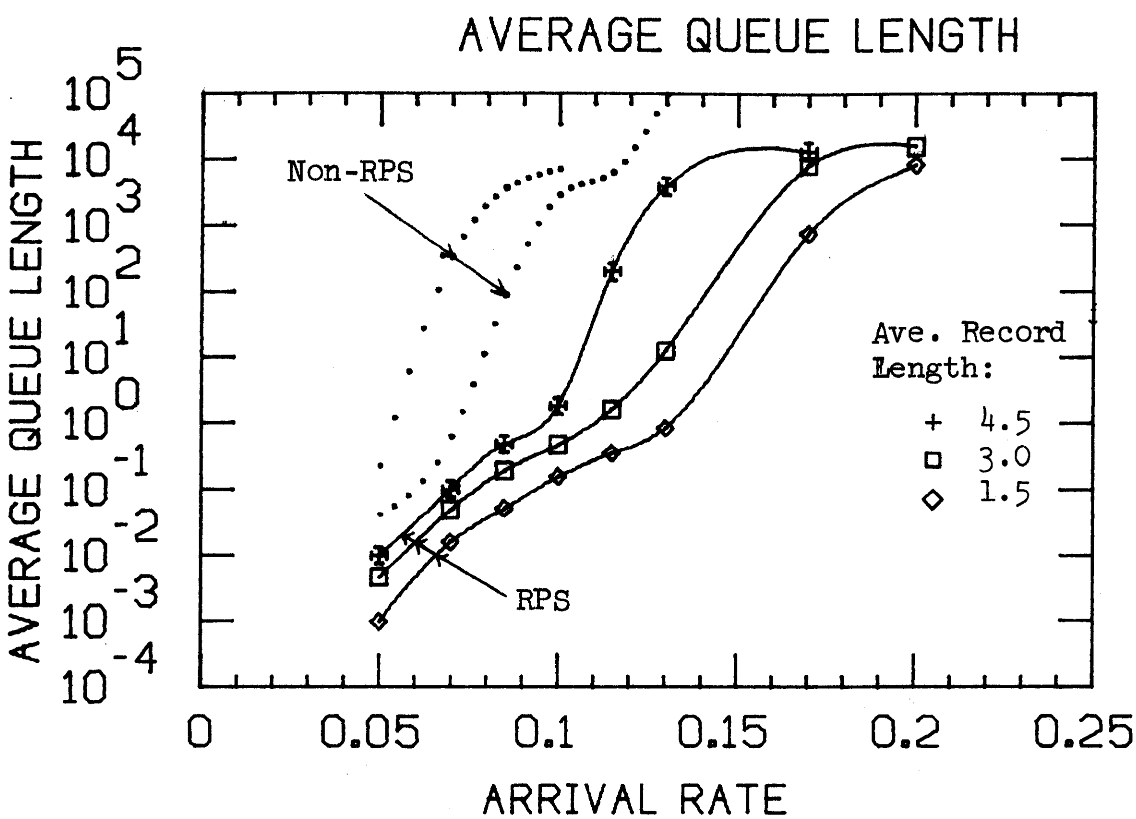
AVERAGE QUEUE LENGTH
AVERAGE QUEUE LENGTH
ARRIVAL RATE
Non-RPS
RPS
Ave. Record
Length:
+ 4.5
□ 3.0
◇ 1.5
10^5
10^4
10^3
10^2
10^1
10^0
10^-1
10^-2
10^-3
10^-4
0
0.05
0.1
0.15
0.2
0.25

Figure 4

SECTION II

PERFORMANCE OF MULTIPLE SPINDLE DISKS WITH SHARED MULTIPLE CHANNELS (ITEL AND IBM CONFIGURATIONS)

1. Introduction

In this report we study and compare the performance of two different Multiple Spindle/Multiple Channel disk configurations. ITEL and IBM schemes are considered. For the former scheme a general analytical model is developed and the results of the analysis are compared with a simulation model of the latter scheme.

2. Disk Storage Configurations

In a general direct access storage system the disk units can be accessed by several channels. An immediate consequence of such a resource sharing environment is the access interferences which must be resolved by the control and queueing mechanisms.

In an ITEL scheme, if we want to connect a series of 7330 disk to, say, two channels, we may use two 7830 disk storage controls (Figure 1a). Each disk storage control can independently control every disk unit. In this scheme, while a disk unit is transferring data via control unit A, and consequently via channel A, Figure 1a, every other disk unit can use the services of control unit B which is connected to the independent channel B.

An IBM scheme can be more restrictive. A pack of 3330 disk storage units which is controlled by one 3333 disk storage control can be simultaneously communicating with only one channel via 3830 control unit. Therefore, when a disk drive is busy transferring data, the other disk drives in the pack cannot receive I/O commands from any other channel which is also attached to the pack.

However, each disk unit can be independently busy executing previous I/O commands like SEEK and RPS. In effect, in this scheme, one disk drive can block other units of the pack while it is transferring information (Figure 1b).

In order to achieve the degree of independence described for the ITEL configuration in an IBM scheme, it seems, one has to use one 3333 for each 3330 in the system.

In the SLAC Triplex System which uses both ITEL and IBM devices, each disk unit may be accessed by three different processors. Therefore, when a disk unit is used by one computer, it is possible that any other disk unit in the same pack is requested by other processors. In Figure 2, the existing IBM 3330 and ITEL 7830 disk organizations in SLAC are shown.

In the following sections we develop models of Multiple Spindle/Multiple Channel disk storage systems, and investigate the effect of access blocking on the performance of the system under hypothetical work loads.

3. Analytical Model (ITEL)

Let us consider a closed queueing network model which describes a rather general Multiple Spindle/Multiple Channel disk storage system as shown in Figure 3. The model consists of d disk drivers, c channels and one processing station. Each disk unit is connected to every channel. The service time of each disk models the actual independent disk activities like SEEK and RPS. The service in each channel consists of transferring one block of information. We allow the I/O commands to reach the disk units even when the channels are busy. In the actual systems the duration of these commands are very short and earlier study (Wilhelm N.C., 1973) has shown that allowing the commands to reach the devices by command subchannels does not significantly alter the performance of the disk storage system.

The processing unit in our model is the simulation of the processing time between the I/O requests in the actual systems.

This network is a conceptual model which captures the basic structure of the disk configuration which we have referred to it as ITEL scheme.

Let us assume there are N requests (customer) in the system. The solution to this model consists of finding the steady state probabilities of the number of requests at each center. We number the service stations starting from the disk units. Therefore, stations 1,2,...d refer to disk units. Next, we number the channels starting from d + 1 to d + c. Finally, station d + c + 1 indicates the processing station.

A state S in this model is a vector of d + c + 1 elements, i.e.,

$$S = (\underbrace{n_1, n_2, \ldots, n_d}_{\text{Disks}} + \underbrace{n_{d+1}, \ldots, n_{d+c}}_{\text{Channels}}, \underbrace{n_{d+c+1}}_{\text{IO Producer}})$$

Where n_i is the number of requests at center i, and $\Sigma_i\, n_i = N$

Let μ_i be the exponential service rate at center i. Furthermore, let the service discipline at each center be on the first come first serve (FCFS) basis.

The solution to the steady state probabilities has a product form (Baskett, et al 1975):

$$P(S) = Kh_1(n_1)h_2(n_2)\ldots.h_{d+c+1}(n_{d+c+1})$$

Where

$$h_i\ (n_i) = \left(\frac{e_i}{\mu_i}\right)^{n_i} \text{ for } i = 1,2,.., c+d+1$$

The coefficient K is normalizing factor so that the probabilities add up to one.

The term e_i's are the solution for the flow equations of the requests is the system as following:

$$
\begin{cases}
e_j = P_{d+c+1,j}\, e_{d+c+1} = \frac{1}{d}\, e_{d+c+1} & j = 1,2,\ldots d \\
e_{d+j} = \sum_{i=1}^{d} P_{i,d+j}\, e_i = \frac{1}{C} \sum_{i=1}^{d} e_i & j = 1,2,\ldots c \\
e_{d+c+1} = \sum_{i=1}^{c} P_{d+i,\, d+c+1}\, e_{d+i} = \sum_{i=1}^{c} e_{d+i} &
\end{cases}
$$

In the above equations $P_{i,j}$ is the probability of going from station i to j.

We have assumed that when station i can reach station j, j = 1,2...n, then $P_{i,j} = \frac{1}{n}$.

A solution for e_i's can be found by letting $e_{d+c+1} = 1$. We then have:

$e_j = \frac{1}{d}$, j = 1,2,..d

$e_{d+j} = \frac{1}{c}$, j = 1,2,..c

$e_{d+c+1} = 1$

We can now find the steady state probabilities by substituing for e_i's the results obtained above:

$$
\begin{aligned}
P(S) &= K \prod_{i=1}^{d+c+1} \left(\frac{e_i}{\mu_i}\right)^{n_i} \\
&= K \prod_{i=1}^{d} \left(\frac{1}{d\,\mu_i}^{n_i}\right) \prod_{i=1}^{c} \left(\frac{1}{c\mu_{d+i}}\right)^{n_{d+i}} \left(\frac{1}{\mu_{d+c+1}}\right)^{n_{d+c+1}}
\end{aligned}
$$

The total number of states is equal to $\binom{N+d+c}{d+c}$

Once we have the steady state probabilities, we can find different performance measures like device utilizations, etc. However, in this case we are only interested in the flow rate of the system. We define the flow rate as the number of requests which arrive to the processing unit in a unit of time.

4. Simulation Model (IBM Scheme)

The previous section was an attempt to model the ITEL configuration. In modeling one pack of IBM disk units which are connected to, say, two independent channels, we have the problem of one channel blocking the other channel while one disk from the pack is transferring data via the former one. Therefore, in our earlier model we make a modification so that while one channel is transferring data the other channel ceases to function, however, its queue may accept the arriving requests at all times.

5. Numerical Results

The essence of our comparison is based on finding the degree to which the channel blocking can effect the flow rate in comparable systems. We know that the service rate at the channels, i.e., the time to transfer data, is an important factor in the result of our comparison. Therefore, we compute the flow rate values with both the number of disk drives and channel service time as parameters.

For the numerical results, we assume the following relative service rates as the basis, and then we increase the channel service rate with respect to the other stations:

$\mu_i = \mu$	$i = 1,2,\ldots d$	disks
$\mu_{d+j} = 5\mu$	$j = 1,\ldots,c$	channels
$\mu_{d+c+1} = 1.8\mu$		processor

These values are not selected purely arbitrary. They are somewhat related to some actual measurements on an aspect of a disk storage system.

For both models we let $N = 7$ and $c = 2$. The flow rate of the first model can be computed in several ways. We first find the steady state probabilities by actually generating all the states and using the results in Section 3. This can be efficiently and easily done with a program with recursive calls. Once we have the steady state probabilities we can find the flow rate by finding, say, the product of the probability of having at least one job at the processing station (d+c+1) by the service rate at this station, summed over all such states.

In the case of the second model, we have resorted to more costly simulation runs. We can, therefore, measure the flow rate by counting the arrivals to, say, the processing station.

In Figure 4 the flow rate is plotted for different number of disk drives (d). The solid lines give the independent channels model (ITEL), and the dotted lines represent the results of the second model with channel blocking in effect (IBM). The pairs of flow rate curves obtained from the two models are plotted for different values of channel service rates.

In this figure we can see that by increasing the number of disk units the flow rate gradually increases accordingly.

We can also see that as long as the channel service rate is high the effect of channel blocking is not significant. For instance, for $\mu_{ch} = 5\mu_{dsk} = 5\mu$ (the top pair of curves) the flow rate of both models are very much the same. However, as we reduce the speed of the channels the difference between the flow rates of the models becomes wider. The bottom pair of curves is for $\mu_{ch} = \mu_{dsk} = \mu$ This fact is also shown in Figure 5 where we plot the flow rate for different channel service rates and d = number of disks = 4.

The choice of N in our models also effects the numerical values of the flow rates. In the earlier computations we assumed N = 7. In order to see where we stand by choosing this number, we plot the flow rate of the first model for different values of N in Figure 6 (d = 4). The choice of N = 7 is a tradeoff between having a reasonably large N to keep devices busy and small enough to keep the state space within computability range.

6. Conclusion

We studied the effect of channel blocking on the performance of disk storage systems by considering two different installation schemes. We observed that channel blocking can be critical in the performance of the system only if the load on channel is very high. For a given load condition, when we reduced the speed of the channel five times (hence, increasing the channel busy time), the throughput of the system with channel blocking dropped to 38% below the throughput of the system without channel blocking.

In this study, we only addressed ourselves to the performance aspects of the two systems. Other factors like cost, reliability, maintenance, etc., are also important in completing the comparison effort. Since both systems are used in this installation, the data should be readily available to complement this study.

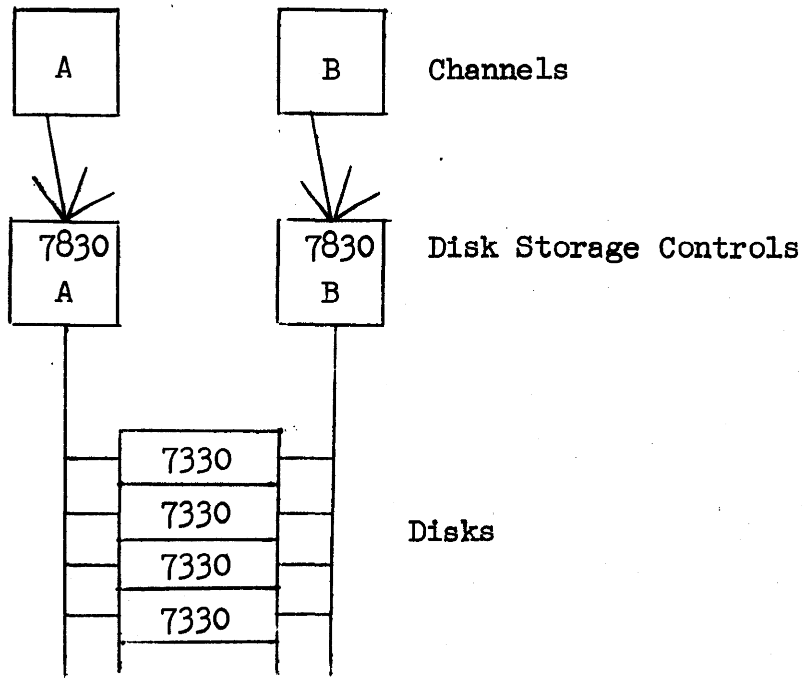

Figure 1a: ITEL multiple channels multiple disks configuration

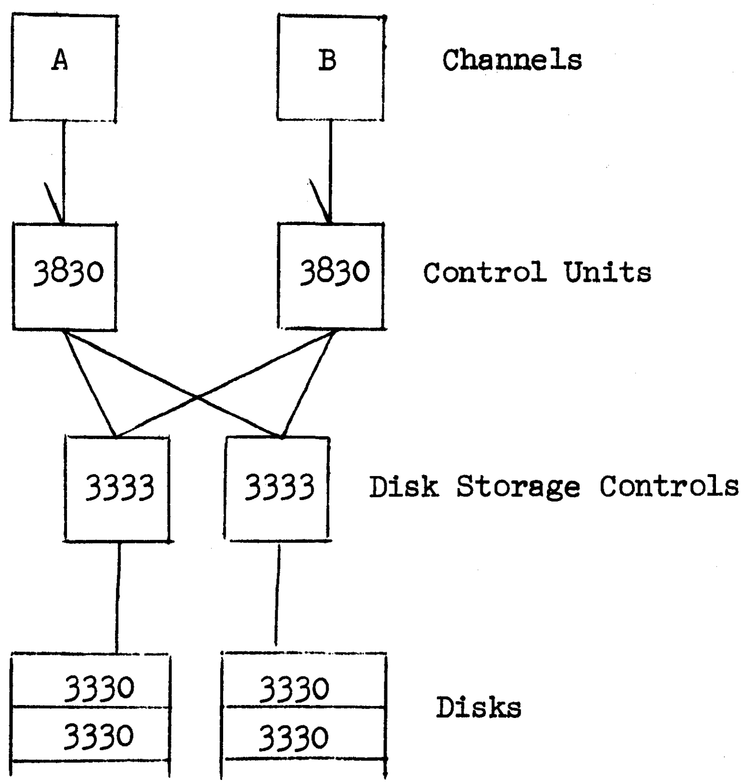

Figure 1b: IBM multiple channels multiple disks configurations

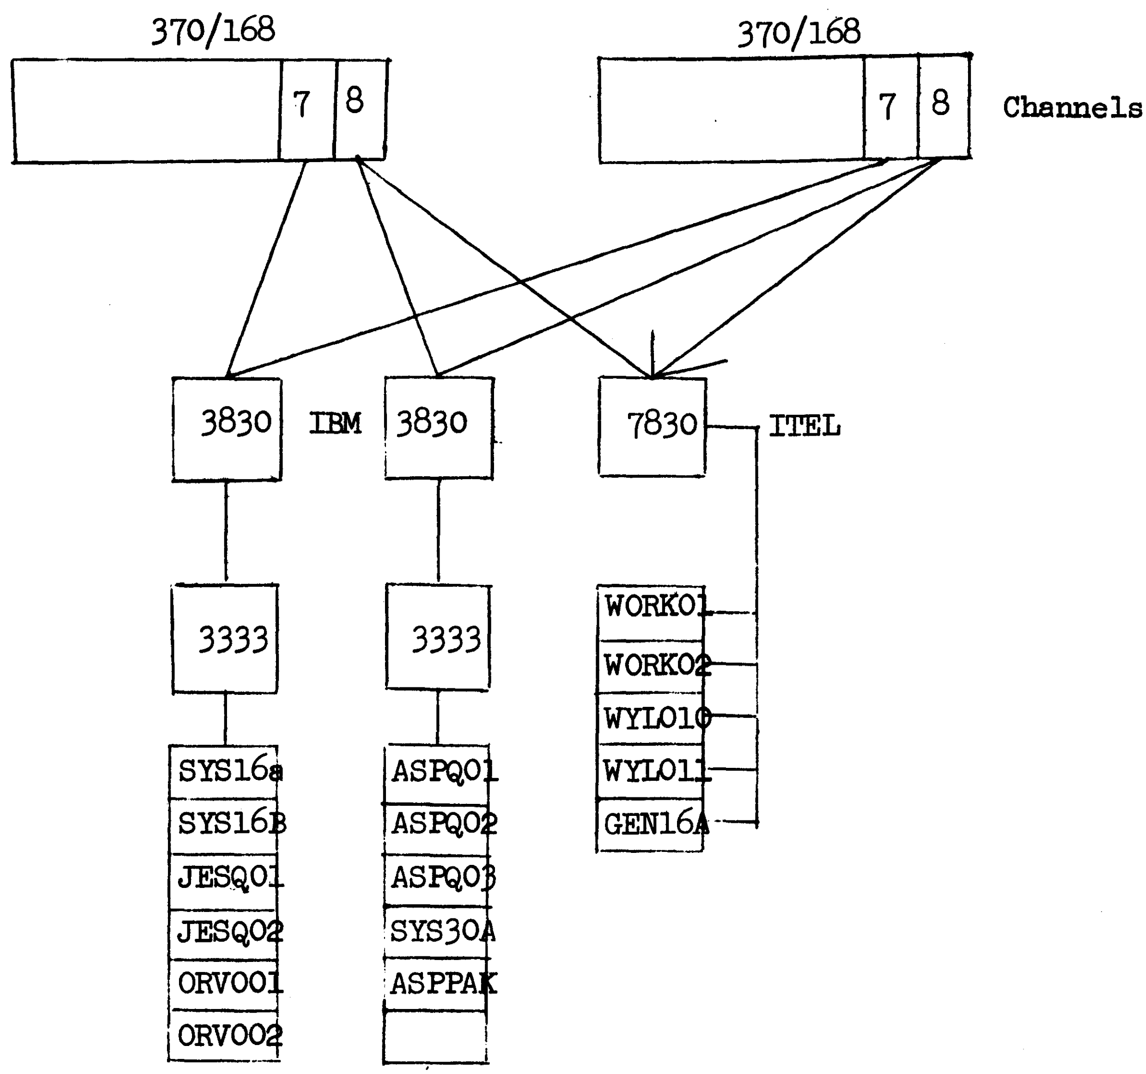

Figure 2: Part of the current disk organization in SLAC Triplex

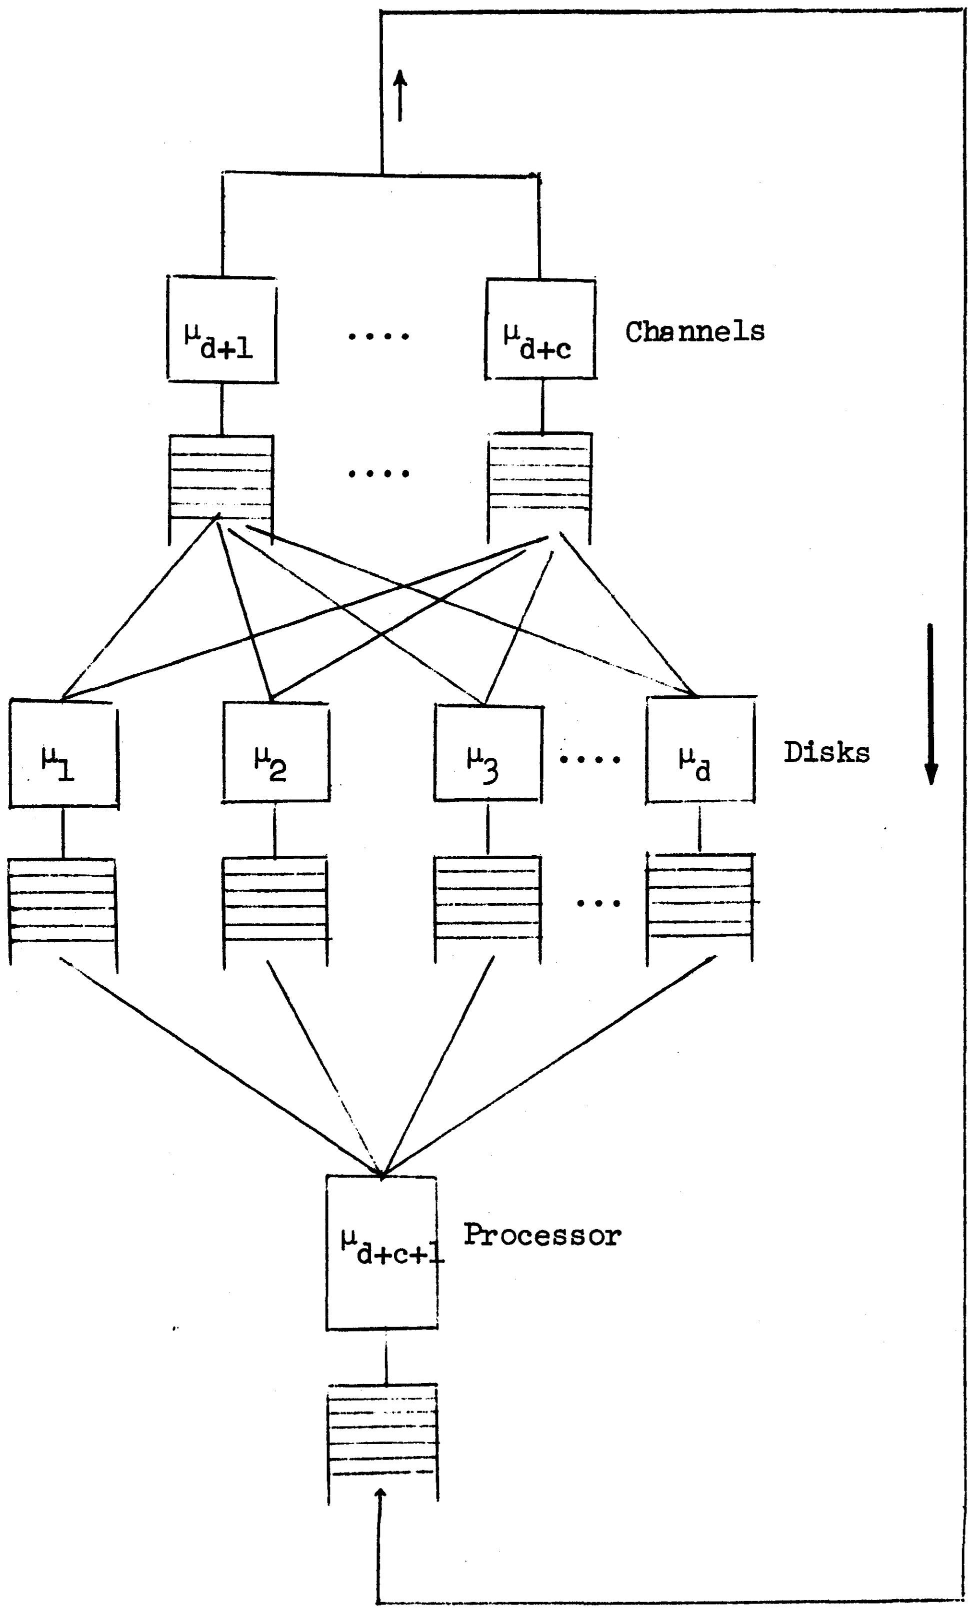

<u>Figure 3:</u> A general Multiple Channels/Multiple Disks queueing model

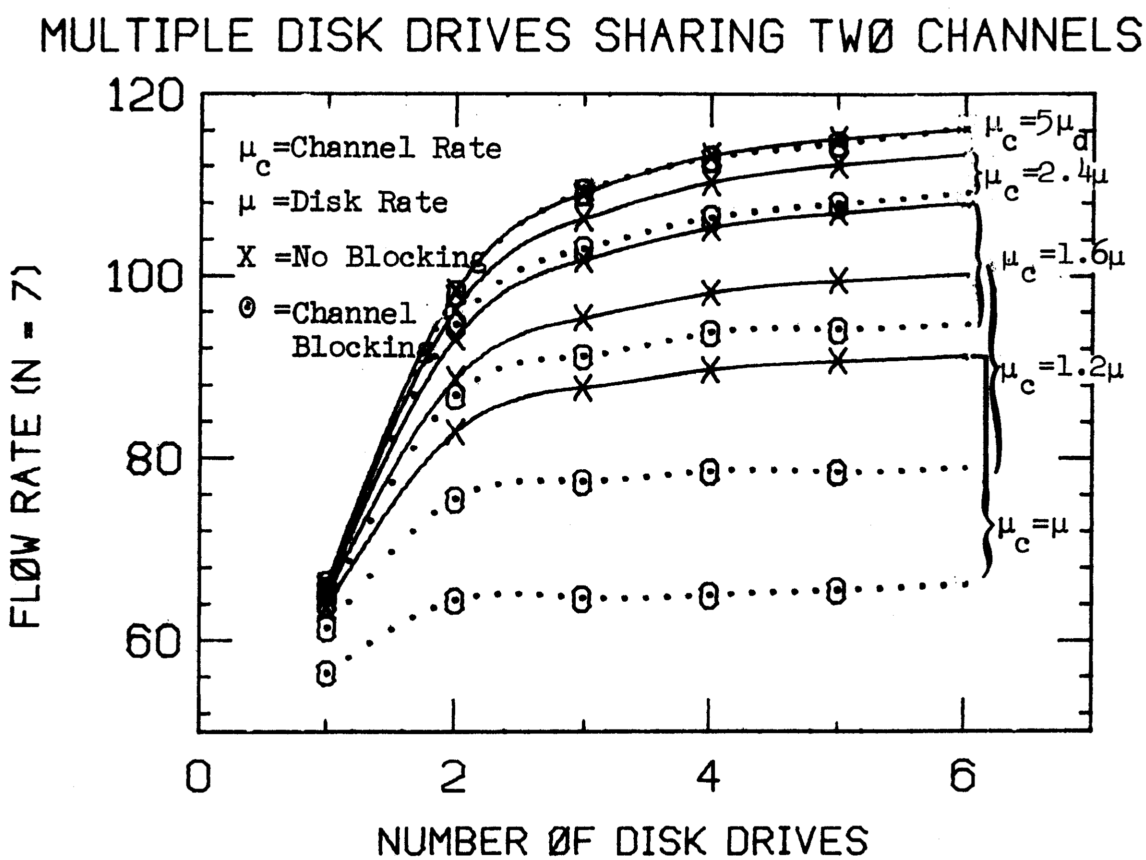

Figure 4 : Comparison of the flow rate in ITEL model with the flow rate in IBM model where channel blocking is in effect.

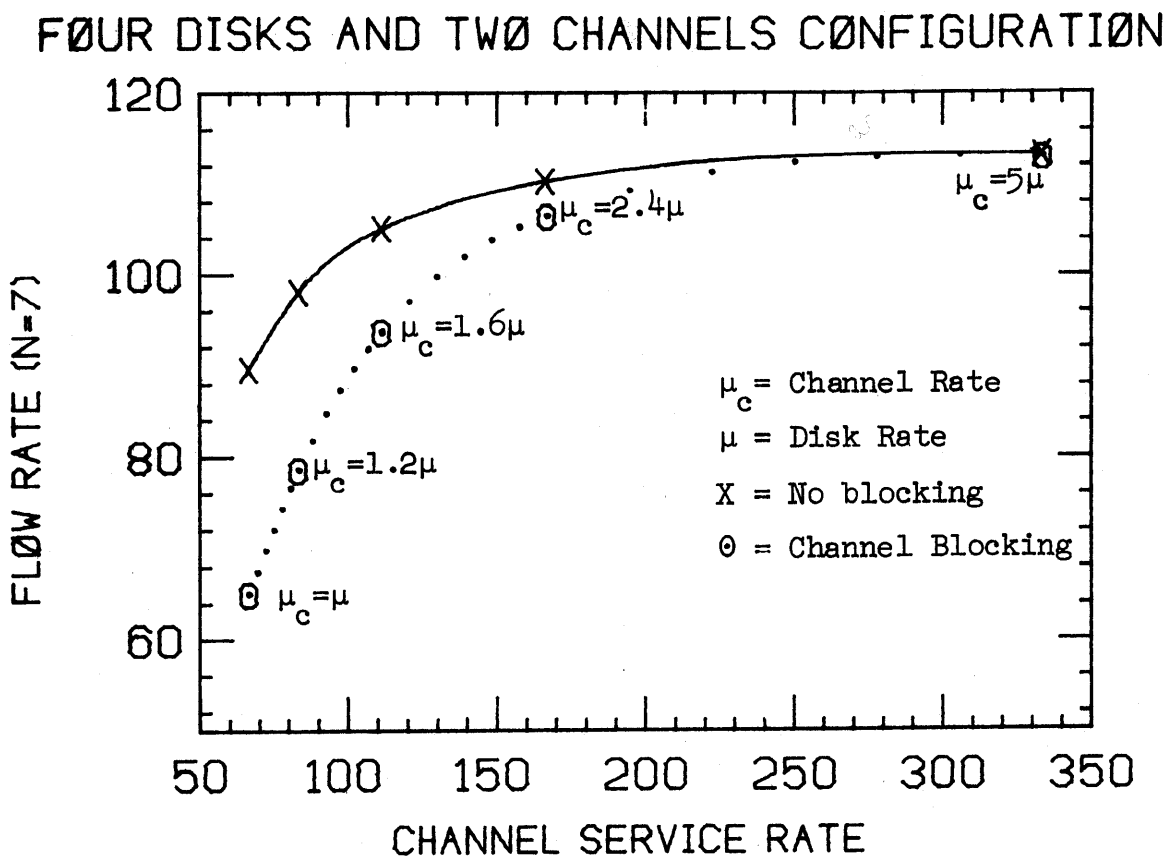

Figure 5: The effect of channel rate in the performance of IBM model compared to the ITEL model where channels do not block each other in transfer time.

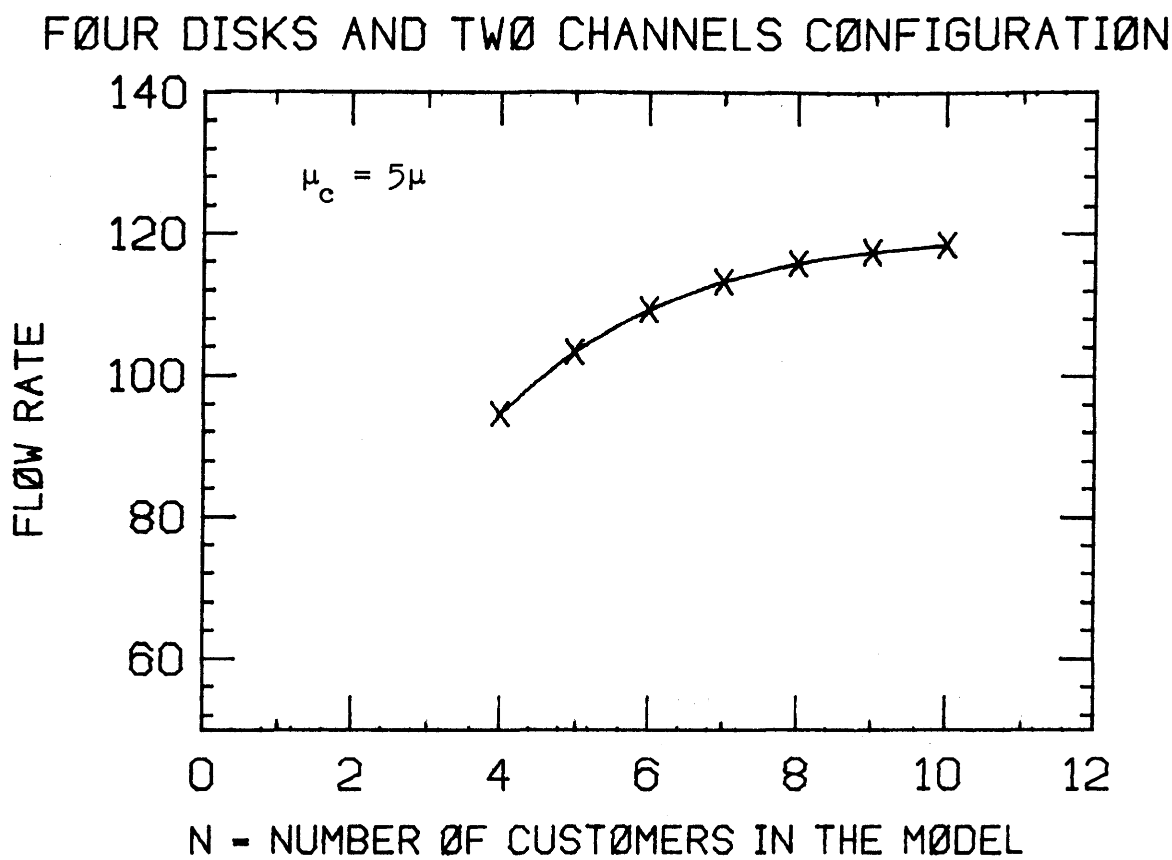

Figure 6: The flow rate of Itel model(no channel blocking) with different number of requests.