663713

DP-1181

AEC RESEARCH AND DEVELOPMENT REPORT

SIMULATION MODEL OF AN IBM 360/65 JOB SHOP - II

M. M. LUDEMAN
F. D. KNIGHT

SRL
RECORD COPY

DU PONT
REG. U. S. PAT. OFF.

Savannah River Laboratory

Aiken, South Carolina

LEGAL NOTICE

This report was prepared as an account of Government sponsored work. Neither the United States, nor the Commission, nor any person acting on behalf of the Commission:

A. Makes any warranty or representation, expressed or implied, with respect to the accuracy, completeness, or usefulness of the information contained in this report, or that the use of any information, apparatus, method, or process disclosed in this report may not infringe privately owned rights; or

B. Assumes any liabilities with respect to the use of, or for damages resulting from the use of any information, apparatus, method, or process disclosed in this report.

As used in the above, "person acting on behalf of the Commission" includes any employee or contractor of the Commission, or employee of such contractor, to the extent that such employee or contractor of the Commission, or employee of such contractor prepares, disseminates, or provides access to, any information pursuant to his employment or contract with the Commission, or his employment with such contractor.

Printed in the United States of America
Available from
Clearinghouse for Federal Scientific and Technical Information
National Bureau of Standards, U. S. Department of Commerce
Springfield, Virginia 22151
Price: Printed Copy $3.00; Microfiche $0.65

663713

DP-1181

Mathematics and Computers
(TID-4500)

SIMULATION MODEL OF AN IBM 360/65 JOB SHOP-II

by

Michael M. Ludeman

and

F. Delano Knight

Approved by

H. C. Honeck, Research Manager
Computer Applications Division

February 1969

E. I. DU PONT DE NEMOURS & COMPANY
SAVANNAH RIVER LABORATORY
AIKEN, S. C. 29801

CONTRACT AT(07-2)-1 WITH THE
UNITED STATES ATOMIC ENERGY COMMISSION

ABSTRACT

A simulation model of an IBM 360/65 computer job shop operating with a dynamic job scheduler is presented. The discussion includes a description of the system modeled, the design of the experiments made with the model, and an analysis of the results obtained. The model is written in GPSS/360.

CONTENTS

LIST OF TABLES AND FIGURES

INTRODUCTION

The use of digital simulation techniques to evaluate the performance of a computer installation has become increasingly useful in the past decade. A simulation study of an IBM 360/65 Computer Job Shop has been reported by Knight and Ludeman.(1) An extension of the model developed in that report is described. Because a familiarity with this earlier report is assumed, many of the details common to both studies have been omitted from this report.

SUMMARY

A series of second-generation simulation models was designed and used to evaluate certain operating strategies in the Savannah River Computer Center in an attempt to improve the operation. The rescheduling of preventive maintenance from 2 days per week beginning at 12:00 noon to 3 days per week beginning at 6:00 a.m. was evaluated with further justification for the rescheduling obtained. Also, a number of priority scheduling algorithms were developed and compared with those in current use. The current HASP priority scheduling algorithm, which considers both expected run time and output in determining priority, was found to be better in most respects than the other algorithms tested. The only algorithms tested which showed improved statistics were those from the group which used only run time in determining priority. The areas of improvement were in the mean job turnaround time in the run time category of 0 to 2 minutes and in the mean overall job turnaround time. These were both a direct result of no jobs being delayed because of excessive expected output.

DISCUSSION

THE SYSTEM MODELED

Equipment and Facilities

The physical system under consideration is similar to the one described in Reference 1; most of the differences have been in the area of operating strategy and environmental modifications rather than equipment changes. The Savannah River Computer Center (CC) has equipment in two locations. The terminal equipment located in the Savannah River Laboratory (SRL) consists of an IBM 2540 Card Reader - Card Punch and an IBM 1403-N1 Printer. (Since all of the equipment considered in the model was manufactured by International Business Machines, subsequent references will be made by model number only, for example, 2540 or 360/65.) All jobs from SRL are submitted and entered into the system at the terminal location, with the exception of the few jobs which must be run on the 7090 Emulator. The Central Processing Unit (CPU) is located in an adjacent building approximately 1000 feet from SRL. The equipment at this Computer Center (CC) includes a 360/65, a 360/30, a 2540 Card Reader - Card Punch, a 1403-2 Printer used by the 360/65, and other peripheral equipment. Also, a second 1403-N1 Printer is attached to the Model 30 and is used for auxiliary printing purposes. One major equipment change was the acquisition of a 2314 Disk Unit, which, in the model, is reserved for system direct access storage space. Three of the original 2311 Disk Units were deleted from the system, and the remaining two were reserved for input/output (I/O) operations.

Description of the HASP System

In the original model,[1] the jobs were processed by what is commonly known as a "batch" processing mode of operation. One alternative to this type of job processing is a sequential mode in which the jobs are processed one after another without any delay between "batches" of jobs. One highly developed way to accomplish this is by means of the HASP (Houston Automatic Spooling Priority) system,[2] which is a specialized program that resides in a single CPU with OS/360. The HASP system has the ability to operate concurrently an almost unlimited number of peripheral devices and do so in conjunction with OS/360 processing. Jobs awaiting various stages of processing are placed in queues on a Direct Access Storage Device, on a dynamic priority basis, thus obtaining the effect of a dynamic job scheduler for all devices controlled by the HASP system. In addition, many jobs can be processed more rapidly because of the extensive blocking and buffering capabilities supplied by the HASP system. To examine the feasibility of using a system such as HASP in the original computer job shop, the original model[1] was converted to a HASP type system by imposing the following restrictions:

1. All operator intervention times, which occurred from the time the job was input into the system until the printout was received, were redefined to be zero.
2. Only one job was allowed to be placed on each job tape.
3. The category definitions were redefined to fit within the scope of the HASP priority algorithm to be subsequently defined.

With these three restrictions, this "Force to HASP" model was run and the results were analyzed. The performance index as defined in Appendix A of Reference 1 was computed to be 0.354 which is a considerable improvement when compared to the best (0.027) obtained for the scheduling algorithms considered in Series III, Run 12, in Reference 1. Run 12 is referred to as the "Original Base Case" in Figure 1 and Table 1.

Job Characteristics

Jobs to be processed originate either at the terminal location in SRL or at the Computer Center between 8:00 a.m. and 4:00 p.m., five days a week, Monday through Friday. These jobs are divided into separate categories, each with clearly defined characteristics. These categories were defined as follows:

Category 1. Express Runs - expected run time less than 6 minutes; expected printed output less than 8000 lines.

Category 2. Regular Runs - expected run time less than 20 minutes; no limit on expected lines of output.

Category 3. Long Runs - expected run time more than 20 minutes.

Category 4. Emulator Jobs - these are sent by the messenger to the Computer Center to be put on tape and run there.

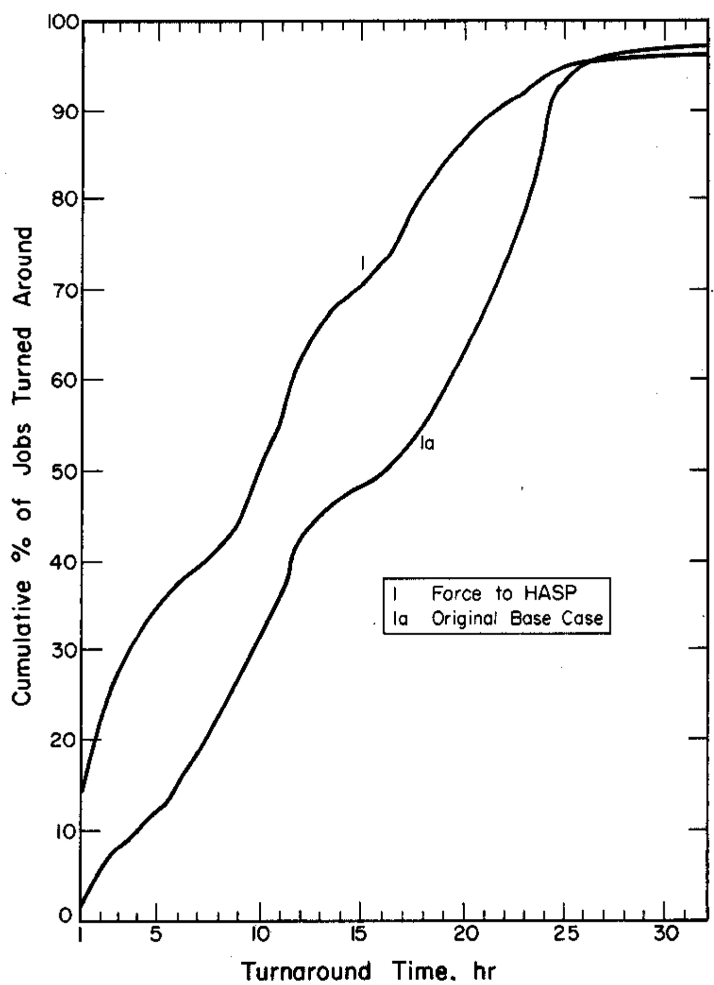

FIG. 1 COMPARISON OF THE ORIGINAL BASE CASE AND FORCE TO HASP MODEL

TABLE 1

OUTPUT STATISTICS FOR FORCE TO HASP MODEL

CPU = 360/65 I/O = 2311 Printer = 1403-2

	Original Base Case	Force To HASP
Job input/day	260	260
Job throughput/day	258	252
Average number of jobs completed, day shift	88	106
Jobs completed day submitted		
Category 1	61	70
2	0	2
Mean adjusted turnaround, minutes		
Category 1	567	359
2	377	373
3	414	423
4	408	402
5	502	481
6	4064	2069
7	453	403
Performance index	0.027	0.354

HASP Priority Algorithm

The order in which the jobs are sequentially queued on the HASP input disk is controlled by what are commonly called the "HASP priority algorithms." These algorithms are generally expressed as a function of expected run time and expected lines of output, although it is entirely possible that such factors as the length of time the job has been waiting for processing could be incorporated into the priority algorithms. For the cases studied in this report, only the expected run time and, in certain instances, the expected lines of output were used to compute this priority. The following matrix (Figure 2) shows the algorithm used when calculating priority based on expected run time and expected lines of output.

Further versatility can be obtained by using one algorithm to schedule jobs during the day shift, while a second is used on the two night shifts.

Expected Output, 1000 lines	Expected Run Time, minutes: 2	4	6	10	30
2	100	90	80	70	60
5	90	80	70	60	50
10	80	70	60	50	40

FIG. 2 CURRENT HASP PRIORITY ALGORITHM

For the sake of uniformity, all priority algorithms will have values ranging from 0 to 110, with priority of 0 to 100 being the normally assigned range and 100 to 110 normally reserved for special systems purposes and jobs that must be completed prior to a deadline.

THE SIMULATION MODEL

Description of the Model

One major change has been incorporated in the model.[1] Earlier it was assumed that jobs which were input from the Computer Center itself were not mixed with jobs input from SRL. They were treated as blocks of time during which the computer was not available to SRL users. In the present model however, an additional card reader and printer have been incorporated and jobs are input from both locations with frequencies representative of the actual input ratio to the Computer Center. A job which is input from either the Computer Center or SRL is returned via the printer to the same location from which it was submitted.

Jobs are generated between 8:00 a.m. and 4:00 p.m. at a rate comparable to a typical input stream. Each job is assigned certain characteristics, which are retained throughout the simulation. These characteristics are assigned to the jobs in a random manner from a distribution that represents the actual characteristics of jobs. These characteristics include:

1. Number of cards in the program deck
2. Processing time on the computer
3. Expected lines of output
4. Number of private tape volumes required
5. Security classification
6. Amount of clerk handling required

The following description of job flow will be in terms of a job submitted from SRL, although it is applicable to jobs submitted from either location. A simplified system flow chart is shown in Figure 3. The jobs arrive at the terminal in SRL and are placed in one of the four categories. If the job is placed in the Express Category, it is immediately processed by the card reader and read into the system. Once the job is read into the system, the job control card is examined to see if the job requires a setup or private tape volume. If none is required, a priority is calculated using the HASP algorithm and the job is placed in a queue on a 2311 disk, with the one with the highest priority placed at the front of the queue. The jobs are then processed in order of priority from the front of this queue.

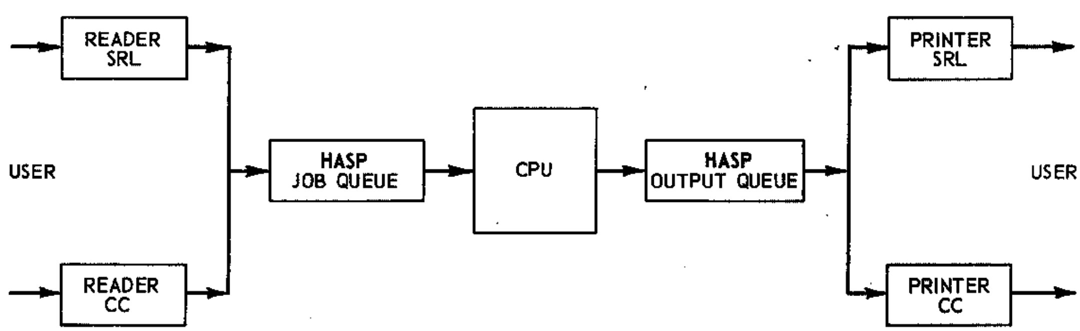

FIG. 3 SYSTEM FLOW CHART - HASP MODEL

If the job requires a special tape, a zero priority is assigned and the operator in the Computer Center physically locates and mounts the tape. This done, the priority is raised to 101 and the job moves to the front of the queue. To prevent the setup jobs from saturating the system, only one job with priority 101 is allowed in the system at a given time.

As each job completes its processing, the front job in the queue finds the computer available and its processing begins. If the job just completed was a setup job, the next setup job in the queue is found; the operator lowers its priority to zero, searches for the private volumes required for it to run, etc. As each job is processed by the computer, its output is queued on a second 2311 disk, again in order of priority, and as the printer in SRL becomes available, the jobs are printed and printouts are returned to the user. Occasionally, when a job is known to have output in excess of 20,000 lines, the output is placed on a tape and then printed via the printer connected to the 360/30. In this way, the output phase of the computer is not tied up by any one job for an excessive length of time. This type of action is taken mainly for short jobs run during the day which will have a large amount of output. Occasionally, because of the importance of a particular job, a high priority (which overrides the system assigned priority) may be assigned to the job. This is not frequently done, and is not considered in the models other than in Run 9 where 5% of the jobs are assigned a priority of 105 to demonstrate the overall degradation of system performance when this feature is abused.

Emulator

The only deviation from the above routine is found in jobs from Category 4 which are the Emulator jobs. These jobs are carried to the Computer Center by the messenger each day and placed on a tape. When no other jobs remain in the system, the operators close down the operating system and bring up the 7090 Emulator. When all of these jobs are complete, the Emulator is closed down and the operating system is reinitialized.

INTERRUPTIONS IN THE OPERATION OF THE MODELED SYSTEM

In the original model, each Tuesday and Thursday at approximately 12:00 noon, all equipment was stopped after completing the job then being processed. The system would stay down for 2.5 ±0.5 hours and then return to operation. This represented the time that the IBM Engineers required to perform routine maintenance on the system. In the present model, with the exception of early runs in which this schedule for preventive maintenance was compared with the new schedule, all of the runs were made using the new schedule which considered preventive maintenance on Monday, Wednesday, and Friday, beginning at 6:00 a.m. and with a length of approximately 2.5 ±0.5 hours. Other interruptions in the operation of the system included the initial program load and unscheduled downtime. In the current model, all of the unscheduled downtime was lumped into one category. This was presumed to occur seven times during the 4-week simulated period, and on each occasion the system was down for approximately 3 ±2 hours.

OTHER CONSIDERATIONS IN THE MODEL

Job Input Control

If the jobs do not get through the system within a certain period of time, the input rate is assumed to decrease considerably because there are only a limited number of jobs available from the users.

Much of the job throughput for a given day consists of jobs in the debugging stage which are submitted several times during the day. If one of these jobs does not return, the user does not have an opportunity to resubmit it that same day. A built-in function automatically decreased the rate of input when the number of jobs within the system exceeded 200. Should as many as 400 jobs be within the system, the input rate was reduced to zero. When the number of jobs within the system decreased to 150, the job input rate was returned to normal.

Relation Between Input and Throughput Rates

The job input rate for each run was controlled by permitting only a predetermined percent of the jobs generated to enter the system, with the remainder terminating. Because this was done by the program on a probabilistic basis, a job input rate of 280 jobs per day might result in an actual input of 275 jobs one day and 290 the next, with an average of 280 jobs per day for the entire run. Under ideal conditions, the job input rate and the job throughput rate should be equal. However, if the system becomes saturated, the job input control portion of the model would decrease the input rate and the job throughput for the model configuration would be decreased correspondingly. In general when the nominal job input rate is higher than the job throughput rate, it may be assumed that the system was saturated for portions of the run. Although the original model used a nonuniform function for the mean job interarrival times, it was decided to use a uniform function in the current models for jobs arriving either at the CC or SRL. All runs for these models used a job input rate of 280 jobs/day with 78% originating in SRL.

VALIDATION OF THE MODEL AND MODEL PARAMETERS

Many of the parameters used in the model were obtained by the analysis of data collected by the Computer Center as described in Reference 1. Certain statistics were not readily obtainable or were not believed to be significant enough to warrant further study. In either case, the statistics were estimated by persons familiar with the area and the estimate was used in the model.

EXPERIMENTAL DESIGN AND RESULTS

To make effective use of the models discussed in this report, a series of experiments was designed to gather statistics for comparison between various operational strategies and priority algorithms. These are defined in Table 2.

TABLE 2

DESCRIPTION OF THE NINE RUNS STUDIED BY THE MODEL

Run	Description
1	Force to HASP Model
2	Original HASP Algorithm
3	First-In, First-Out (FIFO) Algorithm
4	Current HASP Algorithm (Preventive Maintenance 12-2 p.m. Tuesday, Thursday)
5	Current HASP Algorithm (Preventive Maintenance 6-8 a.m. Monday, Wednesday, Friday)
6	Continuous Priority Algorithm: High Priority 0-1 Minute
7	Continuous Priority Algorithm: High Priority 0-2 Minutes
8	Continuous Priority Algorithm: High Priority 0-6 Minutes
9	Current HASP Algorithm with 5% Assigned Priority Jobs

Description of Output Statistics

As a representative example of the statistics obtained using the model, the following statistics were tabulated for each of the runs:

Job Input/Day
Job Throughput/Day
Number of Jobs Returned on Day Submitted
Mean Turnaround Times for Jobs in Following Groupings:
- All jobs
- Jobs with run time 0-6 minutes
- Jobs with run time 0-2 minutes
- Jobs with run time 2-4 minutes
- Jobs with run time 4-6 minutes
- Jobs with run time 6-20 minutes
- Jobs with run time greater than 20 minutes.

Also tabulated was the cumulative percent of jobs turned around as the function of turnaround time.

Series I - FIFO vs. Original Algorithm

This series of runs was designed to compare the classical FIFO queuing situation with the original algorithm used by the HASP system upon its installation at SRL. During the day shift between 8:00 a.m. and 4:00 p.m., jobs were placed on the disk on a short-job-first basis. The only jobs input to the system during this time were jobs less than 10 to 12 minutes as these were the only ones which could possibly be run during the day. Then, between 4:00 p.m. and 8:00 a.m. the following morning, jobs were placed on disk according to a priority algorithm which allowed the longest job in the system to be run first. The purpose of this procedure was to allow the longest jobs to be run during the night and should jobs remain at the start of the following day, they would hopefully be short enough to be run during the early part of the day shift.

Also, if during the day shift, the printer backlog should exceed some predetermined amount, e.g., 20,000 or 30,000 lines, a medium-length job of 10 to 15 minutes was input into the system and run while the printer was "catching up" with the printer backlog. Table 3 and Figure 4 show the results of these two runs.

Series II - Original and Current HASP Algorithms

In this section, the original algorithm described in Series I is compared with the current algorithm used by the HASP system. This current algorithm utilizes both run time and expected output in determining a priority and is shown in Figure 2. Table 4 indicates that although the job throughput and the number of jobs returned on the day submitted was about the same for each of the algorithms, the mean turnaround time for almost all classes of jobs decreased considerably when operating under the current HASP algorithm. The reversal of mean turnaround times for jobs from 6 to 20 minutes and 20 minutes plus is a result of the original priority algorithm running long jobs first at night and then running the medium-length jobs when all long jobs had been run. Figure 5 compares the cumulative job turnaround for the two runs.

TABLE 3

OUTPUT STATISTICS FOR SERIES I: ORIGINAL HASP AND FIFO ALGORITHMS

Statistics		Original HASP Algorithm	FIFO Algorithm
Job input/day		280	280
Job throughput/day		277	258
Jobs returned day submitted		181	65
Mean Turnaround Times, minutes			
Run Time, minutes	Percent of Total		
All Jobs	100	506	1133
0-6	80	178	1070
0-2	58	163	1030
2-4	11	154	1080
4-6	8	332	1070
6-20	13	1729	1030
20+	5	613	1180

TABLE 4

OUTPUT STATISTICS FOR SERIES II: ORIGINAL AND CURRENT HASP ALGORITHMS

Statistics		Original HASP Algorithm	Current HASP Algorithm
Job input/day		280	280
Job throughput/day		277	275
Job returned day submitted		181	175
Mean Turnaround Times, minutes			
Run Time, minutes	Percent of Total		
All Jobs	100	506	400
0-6	80	178	108
0-2	58	163	77
2-4	11	154	131
4-6	8	332	279
6-20	13	1729	954
20+	5	613	1831

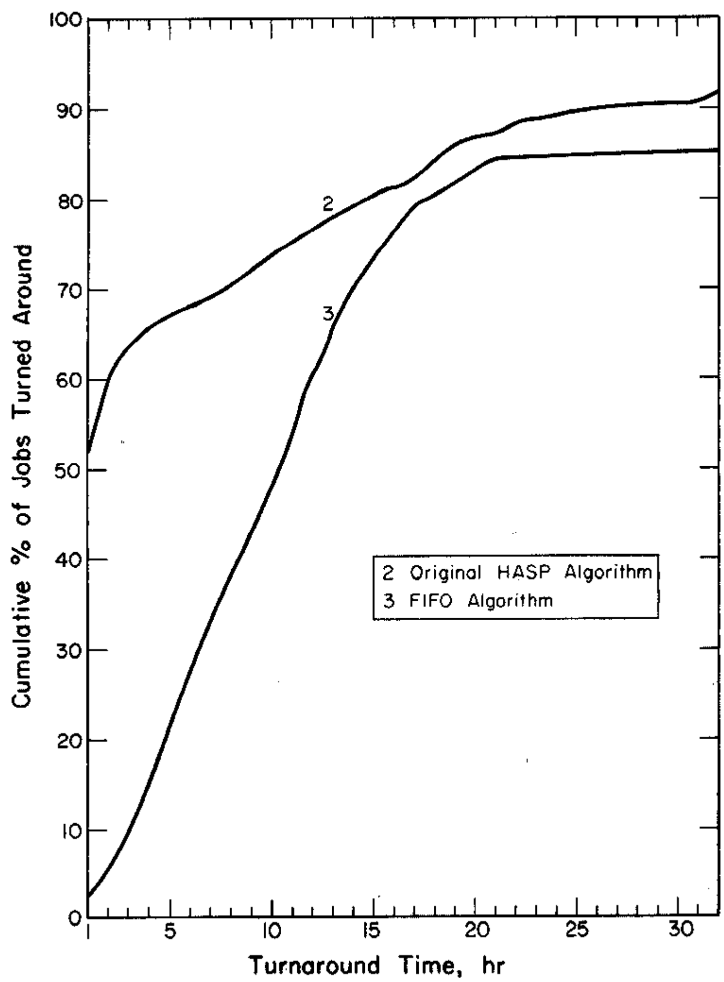

FIG. 4 ORIGINAL HASP AND FIFO ALGORITHMS - SERIES I

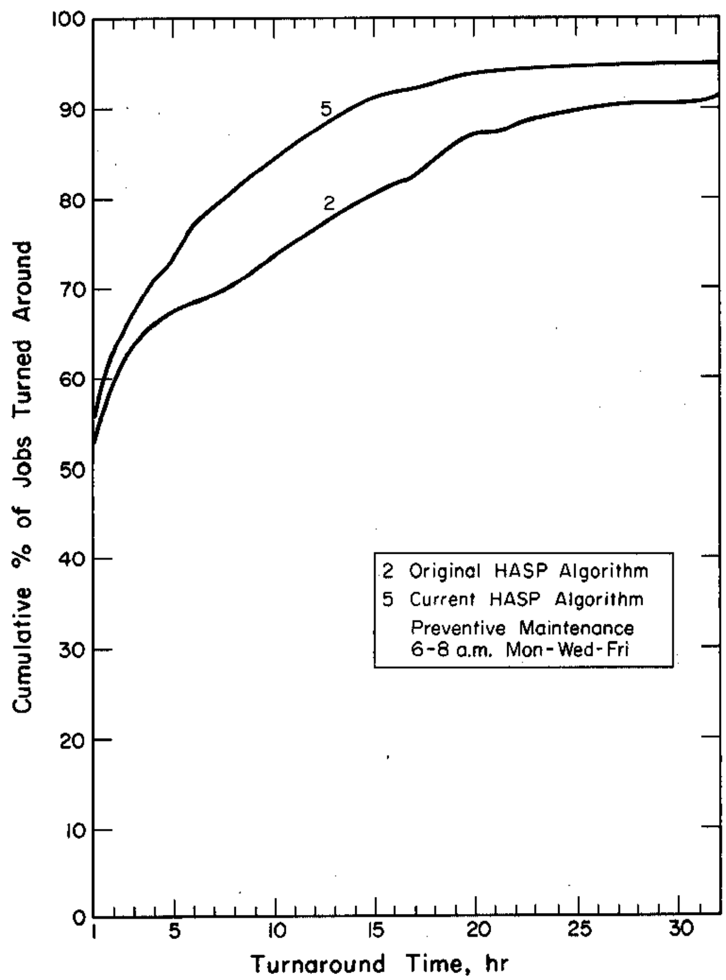

FIG. 5 ORIGINAL AND CURRENT HASP ALGORITHMS - SERIES II

Series III - Rescheduling Preventive Maintenance

In the original model, preventive maintenance occurred twice weekly on Tuesday and Thursday between 12:00 and 2:00 p.m. It had been suggested that a decrease in job turnaround time could be obtained if the IBM Systems personnel were to perform this maintenance at a time other than the day shift. Furthermore, they desired more time to perform this maintenance. Therefore, the preventive maintenance was rescheduled to 6:00 to 8:00 a.m., three days a week, Monday, Wednesday, and Friday. The results of this rescheduling of the preventive maintenance are shown in Table 5 and Figure 6. It can be seen that again, there were no considerable alterations in either the throughput or the number of jobs returned on the day submitted. However, the mean turnaround time for all of the categories was lowered significantly. Both of these runs were made using the current HASP algorithm.

TABLE 5

OUTPUT STATISTICS FOR SERIES III: RESCHEDULING PREVENTIVE MAINTENANCE

Statistics		Preventive Maintenance 12-2 p.m. 2 Occurrences	Preventive Maintenance 6-8 a.m. 3 Occurrences
Job input/day		280	280
Job throughput/day		280	275
Jobs returned day submitted		163	175
Mean Turnaround Times, minutes			
Run Time, minutes	Percent of Total		
All Jobs	100	444	400
0-6	80	177	108
0-2	58	131	76
2-4	11	210	131
4-6	8	420	278
6-20	13	1024	955
20+	5	1189	1831

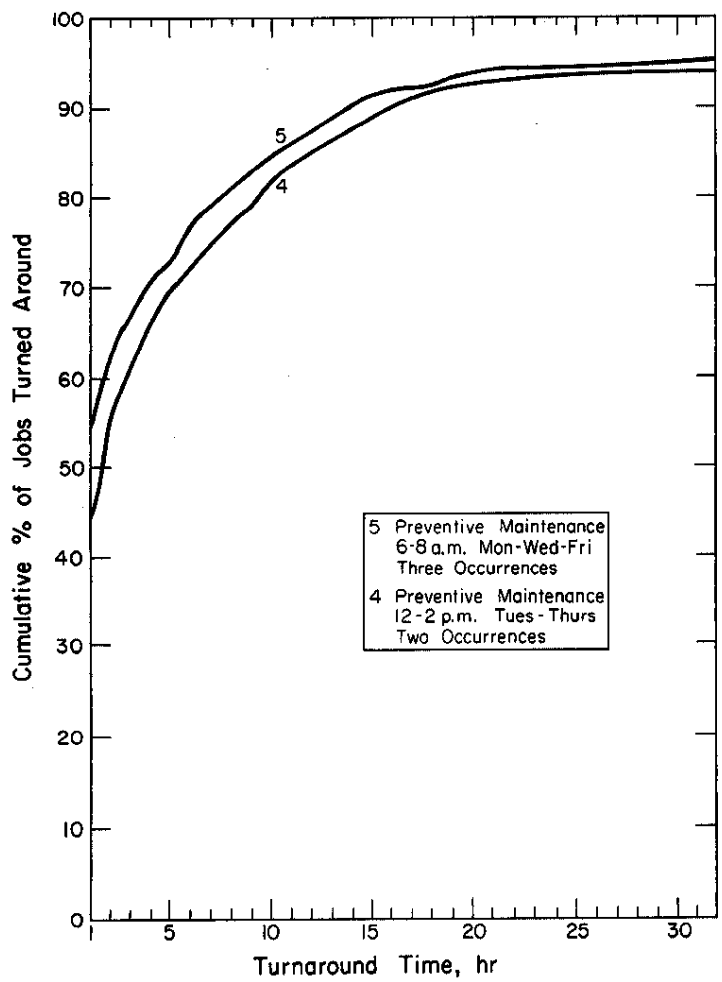

FIG. 6 EFFECT OF RESCHEDULING PREVENTIVE MAINTENANCE - SERIES III

Series IV - Continuous Priority Algorithms

In an effort to obtain a priority algorithm that would minimize all job turnaround, or particularly, minimize the turnaround for jobs in the Express Category, a series of runs was made with various short-job-first algorithms in which the HASP priority was computed strictly as a function of run time. In this series of runs, the algorithms were as follows:

Run 6	High Priority	0-1 Minute
Run 7	High Priority	0-2 Minutes
Run 8	High Priority	0-6 Minutes

These results are tabulated in Table 6 and it can be seen that Run 7 has a slightly lower turnaround time than Run 6 in most of the category groupings shown. For Run 8, however, the run times for 0-6, 0-2, 2-4, and 4-6 are all very close, which would be expected, since all jobs with run time between 0 and 6 minutes were defined to have this highest priority and run prior to any other jobs. By comparing Figure 7 and Figure 4, it can be seen that Run 8 approaches the type of results obtained using the FIFO, queuing discipline for the HASP priority algorithm, which would be expected. As the upper limit on the run time definition for the highest priority class becomes very large, this type of algorithm approaches the FIFO algorithm. Also, by comparing the results from the current HASP algorithm and results from Run 7, it can be seen that the algorithm used in Run 7 provides only a slight improved turnaround time over that of the current HASP algorithm.

TABLE 6

OUTPUT STATISTICS FOR SERIES IV: CONTINUOUS PRIORITY ALGORITHMS

Statistics		Run 6	Run 7	Run 8
Job input/day		280	280	280
Job throughput/day		276	280	279
Job returned day submitted		167	176	108
Mean Turnaround Times, minutes				
Run Time, minutes	Percent of Total			
All Jobs	100	351	339	497
0-6	80	127	114	293
0-2	58	80	63	295
2-4	11	230	180	277
4-6	8	370	360	290
6-20	13	662	684	704
20+	5	1770	1837	1980

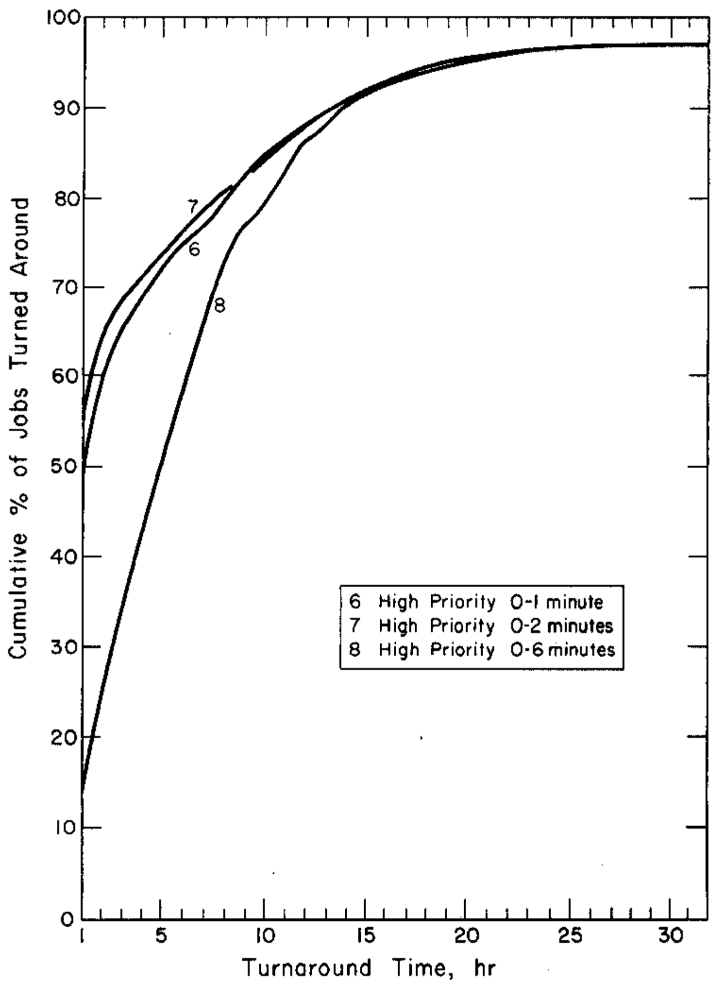

FIG. 7 CONTINUOUS PRIORITY ALGORITHMS - SERIES IV

Series V - Effect of Assigned Priority Jobs

When operating under the HASP system, it is possible to manually override the priority assigned to a job by HASP and assign the job a new and higher priority. While this feature allows for complete flexibility in scheduling, any abuse in its application can cause a serious degradation in system performance.

To demonstrate this degradation of system performance, the model was run with the assumption that 5% of all jobs were manually assigned a priority of 105. In all other respects, the model was identical to the one used in Run 5.

The results shown in Table 7 show an increase in mean turnaround time in nearly all categories, while Figure 8 shows a decrease in the performance curve for the run. The statistics present a strong case against the misuse of the priority feature.

TABLE 7

OUTPUT STATISTICS FOR SERIES V: EFFECT OF PRIORITY JOBS

Statistics		No Priority	Priority
Job input/day		280	280
Job throughput/day		275	275
Jobs returned day submitted		175	171
Mean Turnaround Times, minutes			
Run Time, minutes	Percent of Total		
All Jobs	100	400	475
0-6	80	108	200
0-2	58	76	157
2-4	11	131	223
4-6	8	283	446
6-20	13	955	983
20+	5	1831	1911

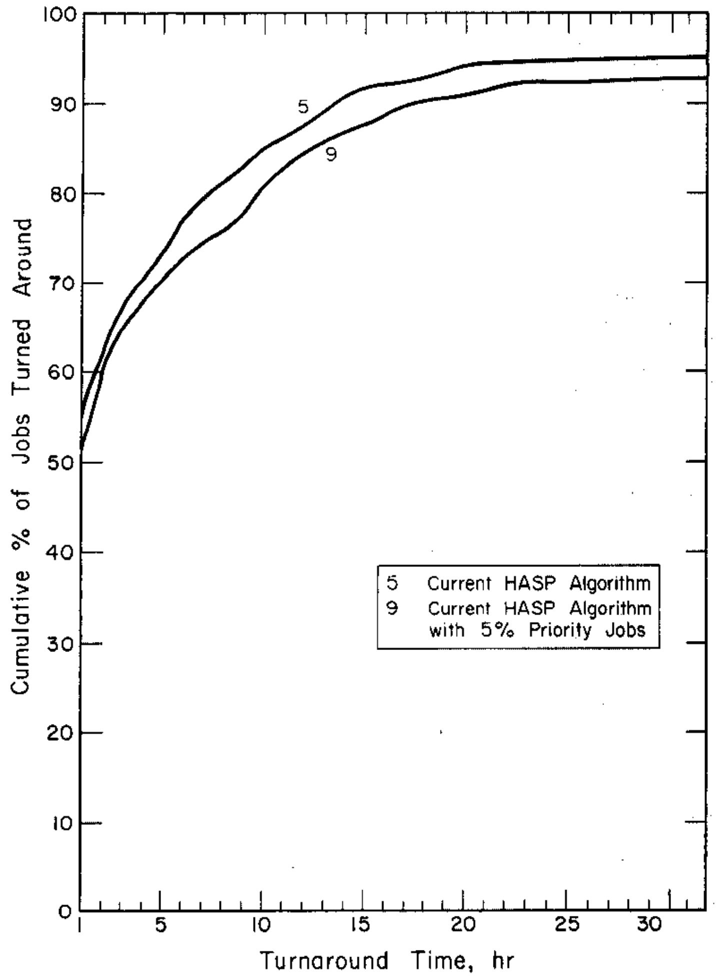

FIG. 8 EFFECT OF PRIORITY JOBS ON THROUGHPUT AND TURNAROUND - SERIES V

APPENDIX A

THE CYCLIC MODEL OF HASP

In the original GPSS model and the HASP models, it was assumed that each job was processed by the system only one time and then terminated. This is rarely the case. A more realistic approach consists of allowing a job to recycle through the model after an appropriate time delay. This delay may be considered to be the time required for a programmer to correct a programming error or perhaps the time required to evaluate computed data before resubmitting the program with new data. This approach was considered to be desirable in an attempt to measure the effect on the system of the very rapid job turnaround time for jobs with 1- and 2-minute run time. The large percent of short jobs run during the 8-hour shift was believed to consist of only a small number of jobs submitted many times during the day. The faster the job was returned to the user, the sooner it would be resubmitted, and the larger the percent of short jobs processed by the system. This led to a vicious circle effect from which escape was practically impossible, save for the end of the work day. With a strict short-job-first scheduling algorithm, very few jobs longer than 3 or 4 minutes would be completed during the day shift.

The difficulty in evaluating this approach was the absence of input data for the job characteristics for this model. The available statistics were equivalent to the job characteristics for the total of all jobs which terminated over the length of the run. In order to obtain the actual input distribution, it was necessary to estimate the job input characteristics and to run the model with a suitable recycling factor and then compare the output job characteristic distribution with that from the Computer Center statistics. By varying the input parameters until a suitable output could be obtained, a good approximation would then be available for the actual input distribution of job characteristics.

Another problem arises in estimating the input frequency for new jobs and the termination frequency for the old jobs. In order to prevent the number of jobs within the model from increasing too rapidly, job termination rates were chosen to be 10%, i.e., if the Computer Center has a total of 150 distinct jobs within its confines then each job terminated with a probability of 0.1 and recycled with a probability of 0.9. To simulate the input of new jobs into the system, a special job generation routine generated 15 new jobs per day. Because of these difficulties in obtaining job input characteristics, this model was never used to compare scheduling algorithms, although several test cases were examined to verify the logic of the program.

REFERENCES

1. F. D. Knight and M. M. Ludeman. *Simulation Model of an IBM 360/65 Job Shop - I*. USAEC Report DP-1149, E. I. du Pont de Nemours and Co., Savannah River Laboratory, Aiken, S. C. (1968).

2. T. H. Simpson, et al. "Houston Automatic Priority System." *IBM Contributed Program Library*. 360 D 05.1.007, August 1967.