VM/370 OPERATING STATISTICS

	JULY	AUGUST
Number of Sessions (excluding MVT/TSO)	10,062	12,415
Connect Hours (excluding MVT/TSO)	10,578	13,844
CPU Hours (excluding MVT/TSO)	265	325
Active Projects	478	496
Active Programmers	589	648

TURNAROUND TIMES FOR OS/MVT JOBS (MIT USAGE ONLY)

This chart shows the turnaround times of MVT jobs run by MIT users during July and August. SRI WEEKLY jobs are not included in the chart because their turnaround times are not meaningful. The chart shows the number of jobs run, the mean (average) turnaround time and the time in which 65% of jobs were returned. All times given are in minutes.

	(JULY 1978)									(AUGUST 1978)								
SRI:	LOW & DEFER			STANDARD			HIGH			LOW & DEFER			STANDARD			HIGH		
	#	AVG	65%	#	AVG	65%	#	AVG	65%	#	AVG	65%	#	AVG	65%	#	AVG	65%
All	4084	177	184	6160	51	31	1500	24	13	5374	97	48	6889	28	12	1686	22	10
Setups	986	263	332	1099	104	93	586	29	24	1575	149	111	1307	55	40	561	26	20
Other	3098	150	131	5061	39	24	914	20	7	3799	75	32	5582	22	8	1125	20	5
Time Submitted*																		
Weekdays:																		
8-10A	361	113	13	420	31	10	202	19	11	475	57	12	528	13	5	263	20	13
10A-12N	446	336	498	1024	70	59	245	22	19	674	164	110	1130	25	17	246	16	10
12N-2P	384	370	421	842	77	43	190	24	16	480	160	214	973	44	24	225	22	10
2-4P	546	289	312	999	76	60	210	38	17	689	171	206	1148	39	27	340	21	17
4-6P	590	256	252	758	53	42	164	35	14	729	150	129	924	31	18	224	9	8
6-8P	306	128	148	424	17	6	115	12	9	403	56	31	505	6	4	103	4	4
8-10P	355	43	22	504	32	6	84	6	5	472	20	12	533	7	3	50	27	6
Weekends:																		
9A-5P	506	10	5	528	11	3	133	24	12	550	6	4	494	14	1	93	45	13

*Note that figures shown here exclude jobs submitted during certain night and evening hours, and therefore do not sum to "All". Times given exclude the time required for physical output, since OS/MVT files are "spooled" to VM for printing.

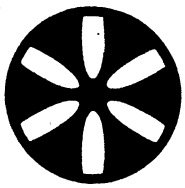

===== VM/370 USER FORUM =====

Monday, October 23
4:00 PM
Bush Room (10-105)

DYNAMO ON MULTICS

by William A. Shaffer

The author is a Senior Consultant for Pugh-Roberts Associates in Cambridge.

A DYNAMO compiler is now available on Multics. Current DYNAMO users will find that Multics DYNAMO fills an important gap, and Multics users may decide that DYNAMO offers a convenient and efficient way of developing computer models.

WHAT IS DYNAMO?

DYNAMO is a language developed in 1959 by the MIT Industrial Dynamics Group to assist in the modeling of dynamic systems. This type of modeling--sometimes called "continuous simulation"--views a system as a network of stocks (levels) connected by flows (rates), in which information about the stocks may affect the flows in various ways. Originally, the Industrial Dynamics Group (now the System Dynamics Group) used DYNAMO to analyze the dynamic behavior of inventory and ordering relationships and of similar industrial systems. DYNAMO's uses soon expanded into branches of the social and biological sciences. Forrester's Urban Dynamics (Cambridge: MIT Press, 1970) and Meadows' The Limits to Growth (New York: Universe Books, 1972) describe two such applications. A considerable (albeit lesser-known) range of applications also extends into the fields of corporate management and health care.

Although the largest group of DYNAMO users comprises researchers involved in System Dynamics, the language can be successfully applied in other fields. Viewed in one way, DYNAMO simply solves a set of simultaneous, nonlinear, first-order differential equations. It is thus useful for simulating almost any such system of equations.

The original developers (especially Professor Jay W. Forrester and Alexander L. Pugh) felt strongly that computer programs should serve their users and not vice versa. In particular, they wanted a language that could be easily explained to people who might lack computer expertise. Their result, DYNAMO, has the following features:

(1) It is easily learned. In courses at MIT's Sloan School of Management, DYNAMO is taught to students at all levels in as little as one or two days.

(2) It is easily debugged. DYNAMO performs extensive error checking and reports errors in English messages. Execution errors (such as division by zero) are reported with the variable name and the equation in which the error occurred, eliminating the need for tracking down such errors with dumps or other tools. Moreover, whole classes of errors commonly encountered in general-purpose languages are either impossible in DYNAMO or are reported much earlier in the debugging process; as a result, programmer productivity is increased.

(3) It eases the task of documenting a computer model. Because the DYNAMO compiler resequences model equations into appropriate computational order, you can place equations in the order that best suits your documentation needs.

(4) It is highly transportable. In the past few years, DYNAMO compilers have been developed for a wide range of computers from fairly small minicomputers like the PDP-11V03 to the major large mainframes. Many universities, government agencies, private organizations, and service bureaus run DYNAMO.

WHY USE DYNAMO ON MULTICS?

DYNAMO is now supported on three computers at MIT. Each system supports different DYNAMO compilers and offers different tradeoffs of cost and convenience. The table on the facing page describes the various DYNAMOs.

As the table shows, the IBM 370 provides the widest variety of systems and the most powerful software. The DYNAMO III/370 compiler offers the most advanced DYNAMO features available. However, students have often found that course and thesis budgets cannot purchase enough computer time on the IBM machine. As a

COMPUTER	PROGRAM	FEATURES	COMMENTS
PR1ME 300 (DOS)	Mini-DYNAMO	basic DYNAMO	inexpensive; limited access
Honeywell 6180 (Multics)	Mini-DYNAMO	basic DYNAMO	inexpensive during Shifts 2, 3, 4; extensive access
IBM 370/168 (CMS)	DYNAMO II/370	basic DYNAMO, plus comparative plots and user-defined macros	
IBM 370/168 (TSO only)*	DYNAMO III/F	all features of DYNAMO II/370, plus arrays	will be phased out as DYNAMO III/370 is perfected
IBM 370/168 (CMS)	DYNAMO III/370	all features of DYNAMO III/F, plus others	experimental version; more efficient than DYNAMO III/F
IBM 370/168 (CMS)	Gaming DYNAMO	capability for writing management games	
IBM 370/168 (CMS)	Documentor	package to aid in documenting DYNAMO models	

result, much student use of DYNAMO has been transferred to the PR1ME 300 at the IPS East Campus Computing Facility. However, while the PR1ME offers inexpensive access to DYNAMO, it supports only a limited number of users, and is available only to students and staff of the MIT Sloan School.

Multics provides a compromise between the IBM 370 and the PR1ME. A large interactive system supporting many terminal types, Multics offers inexpensive usage rates during Shifts 2, 3, and 4. The Multics DYNAMO compiler is Mini-DYNAMO, a language adaptable to a wide range of both small and large computers. The Multics Mini-DYNAMO is identical to the Mini-DYNAMO on the PR1ME. It offers the basic features of DYNAMO, including simple specifications for plots and tabular output, automatic ordering of equations, built-in functions (e.g., delays), and the rerun feature (for repeating simulations with parameter changes). User-defined macros, comparative plots, external functions, and arrays are not supported, but even with these restrictions, Mini-DYNAMO can accommodate most published and student models. Instructors and thesis-bound students who need to use DYNAMO will probably choose to use it on Multics.

GETTING MORE INFORMATION

A good introduction to basic DYNAMO is given in Jay W. Forrester's Principles of Systems (Cambridge: Wright-Allen Press, 1968), available at the Tech Coop.

For a comprehensive discussion of DYNAMO, see the DYNAMO User's Manual. The Mini-DYNAMO User's Guide describes the operation of Mini-DYNAMO. Both are available for purchase in the IPS Publications Office (Room 39-484). An addendum to the latter manual will be available soon and will document the Multics commands used to invoke DYNAMO.

Mini-DYNAMO was developed and is maintained on Multics by Pugh-Roberts Associates, Inc. Questions about applications and operation should be directed to:

Dr. William A. Shaffer
Pugh-Roberts Associates
5 Lee Street
Cambridge MA 02139
Telephone: (617) 864-8880

*TSO accounts are no longer available.

IML RUNOFF ENHANCEMENTS

IPS's version of the "runoff" text-formatting program on Multics was revised August 22. New features allow you to (1) control hyphenation line by line, and (2) feed certain terminal escape functions (such as half-index) to Multics from runoff output.

Hyphenation is normally enabled by the "-hph" control argument on the runoff command line. You may now further regulate hyphenation from within the ".runoff" input file by inserting .hyn (hyphenation on), .hyf (hyphenation off), and .hy (return to initial hyphenation mode) control lines. A .hyf control line tells runoff to stop hyphenating; this is particularly useful for ensuring that the last word of a paragraph will not be hyphenated. Hyphenation is restored to its initial mode (set to "on" or "off" by the presence or absence of the "-hph" argument in the runoff command) with a .hy control line. The .hy control line allows you to produce two types of output (hyphenated and unhyphenated) from one input file. This is helpful when you require both types--e.g., hyphenated copy for distribution and unhyphenated text as input for the create_wordlist and trim_wordlist ("proofreader") commands. The ".hyn {N}" control line may be used to reset the default "breakable" word part for hyphenated output--i.e., the smallest amount of extra output-line space in which runoff should attempt to hyphenate text.

Some computer terminals recognize sequences of characters as instructions for special terminal functions that system software itself can't (or shouldn't be asked to) transmit; some of the characters in these "escape sequences", if supplied to runoff in the input file, might actually print instead of (or in addition to) performing their intended terminal function, thus interfering with the formatting of text. IML runoff will now accept those two-character escape sequences in which the first character is the ASCII ESC character (octal \033) and the second character is one which would normally print. A set of Multics commands can be used in combination with output from runoff to produce documents containing these escape functions. This can be useful, for example, in formatting documents containing subscripts and superscripts, if you have access to a terminal (e.g., a Diablo 1620) which accepts "half-index" escapes.

The procedure for using escapes is somewhat complicated, can be considerably more expensive than straightforward runoff formatting, and may limit your range of output devices. (Escapes vary from terminal to terminal, and can serve only to confuse printers for which they are not intended.) If you feel your application might benefit from this feature, you should consult your terminal's operating manual to verify that the escapes you require are available, and examine the documentation (Revision 1 of "runoff" in Memo MS-13-7) to learn more about the procedure involved.

For further information and assistance, send Multics mail to Tangard.IPSPubs.

MULTICS OPERATING STATISTICS

	JULY	AUGUST	% CHANGE
System hours up	605	751	24.2
Two CPU hours	531	750	41.3
Mean time between fixes (hours)	75:35	187:42	148.3
CPU charged (hours)	475	579	21.9
Memory units (millions)	16.5	22.7	37.4
Print lines (millions)	7.6	9.0	18.1
Terminal hours	14701	19702	34.0
Logins	10974	13567	23.6
Absentee jobs	1872	1899	1.4
Disk usage (page-months)	139160	169762	20.0
Registered projects	395	405	2.5
Registered users	2116	1721	-18.7

SAS AND BMDP UPDATED

SAS and BMDP have been updated on OS/MVT. SAS 76.6C includes several fixes, affecting the following functions and procedures:

AUTOREG	MDY
CHART	NLIN
DATA	PDS
FORMAT	SUM
FUZZ	

IPS has provided a cataloged procedure named SASBMDP to simplify the execution of BMDP programs from SAS. The default REGION specification is 300K; output from SAS is written to file FT12F001, rather than to FT06F001 (as in the SAS cataloged procedure). You must therefore use a /*FORMAT statement for file FT12F001 to get this output. Use the following job format:

```
//  'name',REGION=300K
      <---Job parameter statements
/*FORMAT PR,DDNAME=FT12F001,CONTROL=SINGLE
//  EXEC  SASBMDP
```

The SAS Introductory Guide is now available in the IPS Publications Office (Room 39-484), and includes complete explanations of several of SAS's most valuable procedures, including SORT, ANOVA, and GLM.

The new BMDP-77 statistical package has been installed. In addition to some corrections and improvements, BMDP-77 offers three new procedures: derivative-free nonlinear regression (PAR), general mixed-model analysis of variance (P3V), and life tables and survival functions (P1L). The BMDP 1977 manual (published last year) is now available in the Publications Office. Descriptions of the updates mentioned here are posted in the Programming Assistance and Information (PAI) Office (Room 39-219). Extra copies are available on request. (Contact Julie Hu at 253-6322).

VM/370 OPERATING STATISTICS

	SEPTEMBER	OCTOBER
CMS:		
Sessions	9739	12266
Connect Hours	11563	13667
CPU Hours	67.18	95.59
Page Reads & Writes (thousands)	46559	54675
Non-Spooled I/O Operations (thousands)	30822	30949
Spooled I/O Operations (thousands)	15305	17337
Tape Mounts	650	677
Disk Mounts	381	315
CMS BATCH:		
Jobs	137	183
CPU Hours	0.21	0.13
Page Reads & Writes (thousands)	13	18
Non-Spooled I/O Operations (thousands)	75	61
Spooled I/O Operations (thousands)	122	70
UNIT-RECORD:		
Cards Read (thousands)	2570	2920
Lines Printed (thousands)	5494	7011
Cards Punched (thousands)	37	78
ACTIVE PROJECTS	443	435
ACTIVE PROGRAMMERS	506	554

FASTER ECHOING OVER TELENET

The default modes for ASCII terminals were changed in September to include "crecho" and "lfecho". At that time, Telenet was asked to modify its network defaults to suppress the extra line-feeds it generated locally. Although the new terminal modes are beneficial to local users, Telenet service was noticeably degraded. (Due to network delays, it could take as long as seven seconds between a keyed RETURN and its echoed line-feed.)

If you are a Telenet user, you can speed things up by forcing an explicit reversion to the old protocol, using the combination of Telenet commands shown below after connecting to Multics. User typing is preceded by an arrow (=>); the last three "@" symbols are Telenet command prompts.

```
          Multics 33.15: MIT, Cambridge, Mass.
          Load = 54.0 out of 85.0 units: users = 54
=>        @
          TELENET
          @    =>    INSERT.LF.ON INPUT
          @    =>    INSERT.LF.ON ECHO
          @    =>    CONT
=>        login  YourPersonID  -modes fulldpx,^lfecho,^crecho,...
```

At this point, Telenet will begin echoing line-feeds locally, and your terminal session will run much faster.

NEW BINDER

A new version of the binder (Multics "bind" command) is now installed under >exl>o (the Multics Experimental Library). The new binder accepts input archives in which more than one component has the "trap-on-first-reference" characteristic--a special mechanism within an object segment, used chiefly to implement private error tables. This allows the binding of more than one error table into the same bound segment. (The previous version of the bind command permits only one private error table in a bound segment.) The new binder also includes some internal design improvements that decrease binding CPU time dramatically--by over a third in most cases.

MULTICS OPERATING STATISTICS

	SEPTEMBER	OCTOBER	% CHANGE
System Hours Up	659	708	7.5
Two CPU Hours	653	680	4.1
Mean Time Between Fixes (hours)	109:46	58:59	-46.3
CPU Charged (hours)	387	521	34.7
Memory Units (millions)	18.8	24.8	31.8
Print Lines (millions)	8.2	8.2	-0.3
Terminal Hours	16696	21542	29.0
Logins	14904	18085	21.3
Absentee Jobs	2062	2274	10.3
Disk Usage (page-months)	157102	180749	15.1
Active Projects	322	338	5.0

POSITIONS OPEN

M.I.T. IS AN
EQUAL OPPORTUNITY/AFFIRMATIVE ACTION
EMPLOYER

PLEASE POST

SYSTEMS PROGRAMMER

Systems Programmer in IPS's IBM Technical Services Support Group, to help support the Administrative Computing System running on an IBM 370/148 with VM, DOS/VS, and VS1 operating systems and applications such as CICS/VS. Work entails installing new releases, system tailoring, problem analysis, applying fixes, writing documentation, and user consulting.

Applicant must be proficient in systems programming on most, if not all, of the above systems, and should have experience in a business data-processing environment. A Bachelor's degree or equivalent combination of education and related experience is required.

SYSTEMS PROGRAMMER

Systems Programmer in IPS's Multics RDMS Support Group. This group is in the process of rewriting its Relational Data Management System. Position involves documentation and maintenance, and helping current users convert to the new system. RDMS is written in PL/I and makes significant use of the Multics virtual memory.

Applicant must be proficient in systems programming with PL/I and must have experience on time-sharing systems. Previous experience in database systems and user consulting is desirable. A Bachelor's degree or equivalent combination of education and related experience is required.

Please send resumes to:

Roger A. Roach
IPS Systems Programming
MIT Room 39-551
Cambridge, MA 02139

SPSS CONVERSATIONAL STATISTICAL SYSTEM (SCSS 0.2): COST COMPARISON AND HINTS

by Julie Hu

The introduction of the new SCSS package prompted us to carry out a small benchmark. SPSS on batch (SPSSX), SPSS on TSO (SPSSTSO), and SCSS were tested in a cost comparison. Two data files were used, one containing 220 cases and 23 variables, the other 1017 cases and 167 variables. Four statistical procedures were applied: data analysis, cross-tabulation, Pearson correlation, and partial correlation. The results were as follows:

	Cost for SCSS Run	Cost for SPSSTSO Run
FILE 1 (avg. $2.43 for batch SPSS run)	$2.48 (+2%)	$ 3.09 (+27%)
FILE 2 (avg. $6.12 for batch SPSS run)	$6.98 (+14%)	$10.47 (+71%)

Some hints for SCSS users follow:

(1) A region size of 200K must be specified to provide enough core for SCSS. The TSO LOGON command line should include SIZE(200) after the user ID. (SCSS is used through TSO.)

(2) Data analyzed by SCSS must be contained in a master file, which may be created with SPSS (Release 7) from raw data or from a system file. No raw data file is currently accepted directly by SCSS.

(3) The procedure for executing an SPSS run which includes a SAVE TRANSPOSE statement (creating an SCSS master file) is identical to that for a run including SORT CASES and SAVE FILE: data definitions are required for FT04F001, SORTLIB, SORTWK01, etc.

(4) Three different styles of prompting are available: NORMAL (default), VERBOSE, or TERSE. For a new user, VERBOSE is recommended.

(5) When modifying a variable with the REVISE command (equivalent to RECODE in SPSS), the lowest and highest value in the value ranges must be specified exactly. The keywords HIGHEST (or HI) and LOWEST (or LO) are also permitted. If, for example, the age range is 1 through 92, the following data transformation will cause an error in value ranges:

```
/REVISE AGE
/VALUE 0(MIS);1-25=1;26-50=2;51-75=3;76-100=4
```

To correct this, change 100 to either HIGHEST or 92.

(6) The maximum storage allocated for CROSSTABS or BREAKDOWN of an active table is limited to approximately 4000 cells. In CROSSTABS and BREAKDOWN procedures, only a single table will be formed after a TABLE prompt; the keyword TO is not permitted in the variable list.

5/370 PERFORMANCE

(MIT usage only)

	APRIL	MAY	JUNE
Jobs	18,856	20,272	19,667
Job Steps	36,554	39,270	38,862
TSO Sessions	7,730	8,352	9,809
TSO Connect Time (hours)	3,988	4,273	4840.44
Total CPU Time (hours)	106.39	108.87	118.32
Batch CPU Time (hours)	87.08	87.57	96.11
TSO CPU Time (hours)	19.31	21.30	22.20
Lines Printed (millions)	25.9	25.8	26.1
Cards Read (millions)	5.9	5.8	6.4
Cards Punched (millions)	.5	.4	.6
Tapes Set Up	2,580	2,678	3,428
Disks Set Up	2,215	2,445	2,070
Projects	631	650	628
Programmers	1,103	1,214	849

TURNAROUND TIMES FOR MIT JOBS

This chart shows the turnaround times of MIT jobs run during May and June. SRI WEEKLY jobs are not included in the chart because their turnaround times are not meaningful. The chart shows the number of jobs run, the mean (average) turnaround time and the time in which 50% of jobs were returned. All times given are in minutes.

	MAY 1977									JUNE 1977								
Job Type	SRI LOW			SRI STANDARD			SRI HIGH			SRI LOW			SRI STANDARD			SRI HIGH		
	#	AVG	50%	#	AVG	50%	#	AVG	50%	#	AVG	50%	#	AVG	50%	#	AVG	50%
All	8162	53	18	10870	19	6	554	18	5	8363	56	18	10113	19	7	409	20	6
Setups	1828	60	27	1673	38	19	105	35	11	1974	71	28	1706	35	17	156	36	12
Other	6334	51	16	9197	16	5	449	14	4	6389	52	15	8407	15	6	253	10	4
Time Submitted*																		
Weekdays:																		
8-10A	713	33	15	797	14	6	36	16	6	685	38	13	823	15	6	19	28	9
10A-12N	996	91	54	1645	17	8	72	9	5	1231	80	38	1793	17	10	76	18	9
12N-2P	922	68	39	1511	14	8	97	13	6	1106	88	39	1701	15	8	88	15	6
2-4P	1231	69	38	1823	18	8	103	16	4	1130	74	51	1980	23	10	88	15	6
4-6P	1096	59	39	1591	20	9	115	12	6	1091	42	24	1275	18	8	60	13	5
6-8P	642	32	8	853	14	4	30	6	3	750	18	7	639	12	4	16	7	3
8-10P	901	12	6	819	8	3	19	4	3	816	20	5	646	11	3	25	37	3
Weekends:																		
9A-5P	779	11	4	909	10	3	54	24	2	618	14	5	494	17	4	18	12	3

*Note that figures shown here exclude jobs submitted during certain night and evening hours, and therefore do not sum to equal "All".

MAKING BACKUP TAPES

by Ree Dawson

The Multics system periodically copies (onto magnetic tape) all user storage residing under "user_dir_dir". This service provides a "backup" copy of user storage, so that files lost or damaged by either the user or the system may often be retrieved. A user may also make Multics backup tapes of all or part of his or her storage. For the normal user actively accessing his or her files, making backup tapes is unnecessary--the system automatically provides this service. However, if several files occupy a significant amount of storage and are accessed somewhat infrequently, with periods of dormancy, then making backup tapes may enable a reduction in overall charges by cutting storage costs. Once a user has made a backup tape, files copied to it may be deleted from on-line storage and retrieved from tape when it becomes necessary to access them later.

The Multics backup system was designed for use by the system, rather than by users, so the format and usage of backup tapes is somewhat restrictive. They differ from standard labeled tapes in two critical ways:

- They are readable only by a Multics system, so that they are exportable only to other Multics sites.

- They cannot be made incrementally; that is, using the backup system to write on a tape destroys any information previously stored on the tape.

One advantage of using backup tapes is that the commands which dump and reload user storage may be written as part of exec_coms or ".absin" files, so that backup and retrieval can be relatively automatic processes.

There are two commands, backup_dump and backup_load, which may be used to dump user storage onto tape and to retrieve it on-line from tape later. These commands were not designed for general use, so they are not documented in the Multics Programmers Manual. They are described in a restricted-distribution Honeywell manual, and their descriptions there are oriented towards systems programmers thoroughly familiar with Multics. Users who believe that making backup tapes may be advantageous for them should contact IPS User Services (Telephone 253-7020) for further information and assistance on the use of these commands. The commands are fairly straightforward, once a user knows their default behavior and understands which subset of the many control arguments are relevant to the user's application, rather than to the backup system.

NEW MULTICS MESSAGE AND MAIL FEATURES

by Richard Scott

Over the past few months, several new features have been added to the Multics message and mail facilities. The most recent of these is automatic notification when you receive mail. If someone sends you Multics mail while you are logged in, and if you have invoked the accept_messages (am) command, the message "You have mail." is printed on the terminal. To avoid having this message printed (for example, when running off print-ready copy at the terminal), use the command "defer_messages" (dm), which also defers messages sent via "send_message" (sm). Immediate printing of both types of messages may be restored later by issuing the immediate_messages (im) command.

Another relatively new feature is the "-call" control argument to the accept_messages command. "accept_messages -call cmdline" (where "cmdline" is a Multics command line) causes "cmdline" to be invoked whenever a message is received, rather than printing the message on the terminal. If "cmdline" contains blanks or semicolons, it must be quoted. "cmdline" is invoked with the following arguments (in order): (1) the number zero (unless you are also using the "-hold" control argument to "accept_messages"), (2) the message sender (PersonID.ProjectID), (3) the date-time of the message, (4) the message itself, and (5) the pathname of the mailbox to which the message was sent (unimportant unless you have several). The following example shows how this feature can be used in combination with the exec_com facility to effect an "answering service" when you are away from the terminal.

First, using one of the Multics editors, create a segment named "answer_message.ec" and put the following lines in it (the numbers at the end of each line are for reference only; don't type them):

```
&command_line off                                                          (1)
&print From &2 &3: &4                                                      (2)
&if [not [equal [strip_entry &q2] [user name]]]                            (3)
&then send_message &q2 "Am not at terminal.  Will reply later."            (4)
&quit                                                                      (5)
```

To activate the service, type:

```
accept_messages -call "ec answer_message"
```

Note that, if your working directory is not the one that contains answer_message.ec, you must either change to that directory first or specify the exec_com's absolute pathname, e.g., >udd>PROJ>JDoe>answer_message. When you receive a message in absentia, the contents of answer_message.ec will do the following:

(1) suppress printing of the exec_com's commands as they are executed;
(2) print the sender's name, the date-time, and the message on your terminal;
(3) check to make sure the message was not from yourself (e.g., from an absentee process), thus avoiding a nasty loop;
(4) send the reply "Am not at terminal. Will reply later." to the person who sent the message; and
(5) exit from the exec_com.

The reply sent by line (4) can, of course, be changed to indicate the reason for your absence, or the time you expect to return. To turn off the service when you return to the terminal, type:

```
accept_messages -call ""
```

For further information on the control arguments available for the commands mentioned above, see the individual descriptions in the MPM Commands and Active Functions or type "help <command name>".

MULTICS HINT

We have been pleased (although not yet overwhelmed) by the response to last issue's request for short tips from Multics users. Keep 'em coming! Hints are welcomed from all levels of the user community, and may be mailed to the IPS Bulletin Editorial Office, Room 39-492, or sent via Multics mail to Bulletin.IPCPubs.

Clearly, unique and/or crafty abbrevs are prime candidates for submissions to this column. But they are certainly not the only topic of general interest to Multics addicts. (Besides, too much talk about abbrevs will make the OS folks jealous.) The following hint, submitted by Edward H. Rice (ERice.WFSO), may interest users who frequently develop and share their own data or programming tools.

Users who wish to share segments (readable data, or programs) with other users may find the following abbreviation useful. It ensures that no user may write the segment, sets access for all users, and turns on the safety switch so that the segment cannot be inadvertently deleted.

```
.ab permit do ". delete_acl &1 ..;set_acl &1 re *.*.*;safety_switch_on &1"
```

(The first period within the quoted command line tells abbrev not to look for further expansion within the line, saving a bit of CPU time. The two periods after the delete_acl command provide a match for the existing ACL of any user or project to the segment.)

The reason for denying oneself permission to write the segment is that it reduces the possibility of accidentally modifying information "out from under" another user--i.e., while he/she is using it. In fact, it would be even safer to remove one's own access from the containing directory (i.e. give oneself status permission only) after giving others access to segments. By doing this, it is impossible to accidentally recompile an object segment others are using.

MULTICS OPERATING SYSTEM STATISTICS

	APRIL	MAY	JUNE	% CHANGE*
System hours up	599	635	733	15.4
Two CPU hours	562	623	724	16.4
MTBF (hours)	33:18	48:50	52:20	7.2
CPU charged (hours)	216	189	246	29.9
Memory units (Million)	11.2	11.2	11.6	4.0
Print lines (Million)	10.5	10.3	10.8	4.8
Terminal hours	14774	15034	16108	7.1
Logins	13145	14341	11986	-16.4
Absentee jobs	2583	2343	2455	4.8
Disk usage (page-months)	108919	108482	128271	18.2
Projects	212	216	228	5.6
Users	1650	1656	1507	-9.0

*May-to-June

LATE ITEM: NEW MULTICS RATES

On September 1, 1977, IPS will begin experimenting with a different type of charging for Multics on shifts 2, 3, and 4, and with the concept of user purchases of Multics "ports" for a flat rate. There will also be a new surcharge schedule for users who access Multics through Telenet. The flat-rate shift charges and port charges will be effective through January 1978 and will then be reevaluated.

During the first shift, users will be charged by usage components as currently specified in the IPS memo, "Rates For Computer Services" (GI-7). On other shifts, charges for interactive usage, including CPU, connect time, memory units, and terminal I/O, will be combined as follows:

Shift 2	$4/hour
Shift 3	$2/hour
Shift 4	$4/hour

All disk storage, I/O Daemon, absentee, tape, and retrieval charges will remain unchanged.

For projects with large computing requirements, IPS will offer the option of purchasing a "port" at a charge of $700 per month. This option provides unlimited interactive usage through one guaranteed-access port. (This will not be a directly-wired terminal or private telephone number, but a priority designation to the login process.) Any number of users may be given access by the port administrator, but they may use it only one at a time. (Users can, of course, arrange access through more than one Multics project.) The monthly charge also includes 100 pages of disk storage. Additional disk storage and non-interactive usage will be charged as currently specified in memo GI-7.

The revised surcharge schedule for using Telenet to access Multics is as follows:

Shift 1	$2.25/hour
Shift 2	$1.25/hour
Shift 3	$1.10/hour
Shift 4	$1.80/hour

This surcharge is in addition to first-shift connect charge and the flat-rate charges of the other shifts.

Multics operation shifts are defined as follows:

SHIFT	HOURS	DAYS
1	11:00-18:00	Monday-Friday
2	09:00-11:00 and 18:00-24:00	Monday-Sunday
3	00:00-09:00	Monday-Sunday
4	11:00-18:00	Saturday-Sunday

Inquiries concerning these changes should be directed to Dick Steinberg, Associate Director of IPS, Telephone 253-7184. The revision of memo GI-7 is forthcoming.

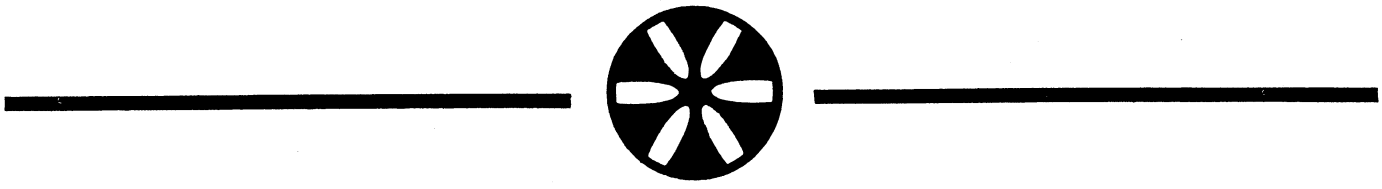

SPSS MAXIVERSION

by Julie Hu

Version 7.00 of SPSSM (Statistical Package for the Social Sciences, Maxiversion) will be installed this month. SPSSM was prepared for users who have been inconvenienced by the 500-variable limitation in the standard version of SPSS (SPSSH); except for capacity, the Maxiversion is identical to the standard version of SPSS. The main differences between SPSSM and SPSSH capacities are:

	SPSSM	SPSSH
variables in system file	1000	500
punched characters in the INPUT FORMAT command	3200	1600
variables in RECODE, COUNT *RECODE, and *COUNT commands	800	400
TRANSPACE = limit for NTRANS*X, NVALS*Y, and NOPERS*Z	X = 164 Y = 41 Z = 20.5	X = 100 Y = 25 Z = 12.5
variables in an archive file	5000	2500

SPSSM should not be used routinely for problems which can be handled by SPSSH, since the larger tables required increase internal overhead, especially for data transformation (as indicated by factors X, Y, and Z above).

Version 7.00 of SPSSM is invoked by specifying the NAME and EXEC statements as follows:

```
     //  'name',REGION=250K
     /*MITID userparm
     //  EXEC   SPSSX,VERSION=M700
or
     //  'name',REGION=250K
     /*MITID userparm
     //  EXEC   SPSS,VERSION=M700[,other symbolic parameters]
```

If you have any questions, please contact Julie Hu in Application Services, Telephone 253-6322.

S/370 OPERATING STATISTICS

(MIT usage only)

Problems in the processor, on a cabling connection in one of the 2870 multiplexor channels, on 3 disk drives and on both 2301 drum storage devices were responsible for the increased number of hardware-related service interruptions on the 370 these past two months. Failures of this type require more than a simple reload. One failure required 21 hours of engineering support before the problem could be identified and corrected--the longest period of down-time we've experienced this year.

	JUNE	JULY	AUGUST
Jobs	19,667	15,905	19,516
Job Steps	38,862	31,303	N/A
TSO Sessions	9,809	7,456	10,189
TSO Connect Time (hours)	4,840	3,316	4,592
Total CPU Time (hours)	118.32	95.19	120.04
Batch CPU Time (hours)	96.11	78.17	96.97
TSO CPU Time (hours)	22.20	17.02	23.07
Projects	628	563	N/A
Programmers	849	775	N/A

TURNAROUND TIMES FOR MIT JOBS

This chart shows the turnaround times of MIT jobs run during July and August. SRI WEEKLY jobs are not included in the chart because their turnaround times are not meaningful. The chart shows the number of jobs run, the mean (average) turnaround time and the time in which 50% of jobs were returned. All times given are in minutes.

	(JULY 1977)									(AUGUST 1977)								
Job Type	SRI LOW			SRI STANDARD			SRI HIGH			SRI LOW			SRI STANDARD			SRI HIGH		
	#	AVG	50%	#	AVG	50%	#	AVG	50%	#	AVG	50%	#	AVG	50%	#	AVG	50%
All	5815	342	25	8938	387	9	443	25	6	7460	62	21	10641	22	8	653	18	8
Setups	1297	103	35	1562	337	21	129	24	13	1813	77	31	1872	36	19	251	23	13
Other	4518	410	21	7376	397	7	313	26	4	5647	57	18	8769	19	7	402	16	5
Time Submitted*																		
Weekdays:																		
8-10A	473	86	31	694	28	11	38	132	7	676	46	21	820	14	7	60	18	8
10A-12N	777	104	71	1521	27	13	98	13	7	965	90	44	1727	18	10	137	14	8
12N-2P	603	652	57	1494	18	10	102	10	6	846	85	40	1742	18	9	91	17	7
2-4P	786	344	71	1753	606	9	82	24	7	1006	81	51	2031	26	11	129	19	12
4-6P	847	1276	36	1271	1502	9	55	8	5	963	68	37	1554	23	9	118	26	8
6-8P	474	25	10	521	10	5	25	19	4	642	22	10	623	12	5	31	11	5
8-10P	647	15	5	392	10	4	23	3	6	843	19	8	577	9	4	37	8	3
Weekends:																		
9A-5P	558	39	7	707	60	6	14	2	2	681	14	4	512	11	3	16	12	5

*Note that figures shown here exclude jobs submitted during certain night and evening hours, and therefore do not sum to equal "All".

delete_acl: Incorrect access to directory containing entry. <dirpath>

and you should contact someone who does have access. Find this information by typing:

list_acl <dirpath>

where <dirpath> is the pathname of the directory containing the damaged segment. For further information on damaged switches, type "help damaged_sw_off" and "help damaged_sw_on".

CHANGES TO IML RUNOFF

MIT's version of the "runoff" text-formatting program on Multics has been revised effective August 8, to incorporate one new feature (marginal change indicators) and important changes to the tabulation control requests.

Frequently, when documents are revised, authors or publishers provide their readers with an easy way to find the differences between the document's previous edition and the current revision. This is usually accomplished by placing a symbol or symbols in the margin adjacent to those sections of text which have been amended. Runoff now provides a facility for automatic insertion of revision indicators ("change markers") in the margin of formatted text. Two control lines, ".bbm" (block-begin for markers) and ".bem" (block-end for markers) are used to enclose altered sections of text in the input file. These controls allow the user to specify the margins to be used for the change markers (left, right, or both) and the character(s) to comprise the change-marker string itself. Different marker and margin patterns may be specified for odd and even pages.

Tabulation in runoff is now implemented by the following control words:

.htd (horizontal tab define)
.htc (horizontal tab cancel)
.htn (horizontal tab on)
.htf (horizontal tab off)

The new controls permit up to 20 different tabulation patterns to be defined and used within a document. Tabulation patterns are assigned a name by the user, and may be turned on and off independently and repeatedly. The user may designate a separate "trigger" character for each tabulation pattern to be replaced by a tab on output, and may specify a character string to fill the "white space" left by a tab. The tabulation control lines used in the previous version of runoff (.sc and .st), while still accepted, are obsolete, and runoff prints a warning message when supplied with input files containing them.

For further details on the use of change-marker and tabulation control lines, see the write-up of the IML runoff command in memo MS-13, now available in the IPS Publications Office, Room 39-484.

MULTICS OPERATING STATISTICS

	JUNE	JULY	AUGUST	% CHANGE*
System hours up	733	600	723	20.5
Two CPU hours	724	580	694	19.7
MTBF (hours)	52:20	49:59	21:54	-56.2
CPU charged (hours)	246	165	190	14.9
Memory units (Millions)	11.6	9.6	11.1	14.8
Print lines (Millions)	10.8	8.2	10.1	23.1
Terminal hours	16108	12538	15257	21.7
Logins	11986	9960	11325	13.7
Absentee jobs	2455	2351	3061	30.2
Disk usage (page-months)	128271	95842	118966	24.1
Projects	228	230	230	0.0
Users	1507	1547	1260	-18.6

*July-to-August

POSITIONS AVAILABLE — PLEASE POST

SYSTEMS PROGRAMMER II
(Level 11)

Systems Programmer II to work on the design, implementation, and documentation changes to operating systems and subsystems. Will also perform system maintenance and crash analysis. Our present system is an IBM 370/168 running OS/MVT with ASP and TSO. The primary assignment is to aid in the conversion to VM/370, CMS, and VS1.

An applicant should have at least 4 years of experience in systems programming with large-scale IBM operating systems. Knowledge of the internals of the current or proposed system or subsystems is desirable. A Bachelor's degree or equivalent combination of experience and formal education is required.

Please send a resume to:

Roger A. Roach
IPS Systems Programming
MIT Room 39-551
Cambridge, MA. 02139
(617) 253-7011

SYSTEMS PROGRAMMER II
(Level 11)

Systems Programmer II to perform system assurance work (crash analysis, installation of changes, performance tuning) and aid in general development tasks (database systems, communications subsystems, writing new commands) for the Multics system.

An applicant must be proficient in systems programming with PL/I and have experience in using Multics or other large-scale time-sharing systems. A Bachelor's degree or equivalent combination of experience and formal education is required.

Please send a resume to:

Roger A. Roach
IPS Systems Programming
MIT Room 39-551
Cambridge, MA. 02139
(617) 253-7011

MIT IS AN EQUAL OPPORTUNITY/AFFIRMATIVE ACTION EMPLOYER

The "search rules" are simply a list of directories maintained for use by the linker. The user can modify this list to add and delete any directories he wishes. The directories are searched in the order they appear in the list. The only "rule" the user should not modify is the one which appears as "initiated_segments". As explained above, the first place the linker automatically checks when looking for a reference name is in the table of previously inititiated segments, and removing this search rule could result in very undesirable behavior on the part of system routines. The initial contents of the list for most users is :

```
inititiated_segments
referencing_dir
working_dir
>system_library_standard
>system_library_unbundled
>system_library_1
>system_library_tools
>system_library_auth_maint
```

Note that the user's working directory is placed before the system libraries by default; if a user has a program with the same name as a system program, the user's program will be substituted for the system program, unless the call is made by another system program in the same directory (referencing_dir) or the system version has been referenced already (initiated_segments).

Comparisons done in 1974 have shown that using the Multics dynamic linkage facility to link and run (one time) a given PL/I program with several external subroutines is more efficient than either loading or linking the same PL/I code on OS/360. This is true because segment names in a directory are hash coded for fastest possible searching, and the segmentation mechanism provides an efficient means for handling the storage allocation problems.

The time required for dynamic linking, in the case of a program which always references the same subroutines, can add up to an appreciable amount when that program is executed many times in different processes. (Note that this is true only for many executions in different processes; subsequent executions of the same program in the same process are direct transfers, as mentioned above.) For these cases (e.g., system programs referencing each other), Multics has a program, known as the "binder", which combines and pre-links programs.

MULTICS OPERATING SYSTEM STATISTICS

	January	February	% Change
System hours up	658	614	-6.7
Two CPU hours	631	452	-28.3
MTBF (hours)	73:06	21:55	-70.0
CPU charged (hours)	183	204	11.2
Memory units (million)	9.6	10.7	11.2
Print lines (million)	10.5	10.4	-0.1
Terminal hours	13738	13069	-4.9
Logins	10918	12001	9.9
Absentee jobs	1917	3449	79.9
Disk usage (page-months)	111449	105787	-5.1
Projects	221	228	3.2
Users	1585	1677	5.8

SURCHARGE ELIMINATED

For several years our batch-job charging structure encouraged users to describe the relative proportion of input or output versus internal computation. The accuracy of these descriptions was important to the effectiveness of the scheduling algorithm in keeping all system components active and in preventing bottlenecks when the IPS IBM 370/158 was heavily loaded for hours or days at a time. Basically, the algorithm attempts to assure that only one CLASS C job is scheduled for MAIN at a time, since a job retains control of the CPU until interrupted by an I/O completion signal.

Now we have a faster central processor (an IBM 370/168), and many of the smaller tasks are performed through TSO, so that the batch-job scheduling algorithm no longer plays such a critical role in system performance. We are therefore, as of March 30, 1977, removing the financial penalty associated with incorrect CLASS specification (At this time, the penalty message will also be reworded to print: JOB CLASS INCORRECT. nn% ADJUSTMENT IN CPU CHARGE RESULTED = $xx.xx). Since there still is system benefit from using this information, we are retaining the discount associated with correctly identifying CLASS=C jobs.

With the new charging structure, a job specified as CLASS=A but whose actual CLASS is C will receive only a warning of incorrect job CLASS, but no surcharge will be added. So if your program goes into an unintended compute loop, you will no longer be financially penalized for the "mis-classification" of your job. Conversely, a job specified as CLASS=C which falls back into the CLASS=A category because it terminated sooner than expected will be charged just enough extra to bring the total up to the CLASS=A rates, but not beyond. It is now safe for a user developing a new program to specify CLASS=C when there is some expectation that it may perform that much computation. It is not advantageous to specify CLASS=C for other jobs, because of possible delays imposed by the scheduling process.

The following table shows the Rate Factor multiplied by the Surcharge Factor to produce the Effective Rate:

Actual Class	Specified Class A-B (0-63% CPU)	C (52-100% CPU)
A-B	1.0 x 1.0 = 1.0	.8 x 1.25 = 1.0
C	1.0 x 1.0 = 1.0	.8 x 1.0 = .8

STILL EXPERIMENTING...

On the first of March, an experiment was started in the PAI Office to discover if our current hours of 9 - 12 noon and 2 - 5 p.m., Monday through Friday, are the most effective hours for users and, thus, the most efficient use of our consultant-power. On that date, we began to keep the office open daily from 9 - 5, with just one consultant in the office from noon to 2 p.m. We also continued to ask the consultants to survey the user load for each hour slot. We are discovering that there is relatively little usage from 9 to 10 (as was expected) and from noon to 1. And the usual 2 p.m. rush is now more spread out!

We will continue these experimental hours and the user load survey until the end of March. Then we plan to use the information which we gain from the survey and any feedback or input from you, along with the benefit of our experience as consultants, to determine that, either our current hours are the most effective or the PAI Office hours need to be rescheduled to reflect the users' needs.

We hope to be able to let you know the results of this effort in the next Bulletin. And, if we feel that a new schedule of hours is necessary, we will try to implement them at the beginning of May. Watch the JPS tailsheet notices and TSO broadcast messages for any changes to the schedule.

We're making a sincere attempt to learn how we can better serve your needs, and hope that this experimentation and survey will be instructive in helping us achieve this goal. But your opinion would be most beneficial, so let us hear it. Contact Debbie Colaw, Room 39-421, Telephone 253-7720.

UPDATES TO P-STAT AND SPSSH

by Suzanne Chen

In SPSSH version 6, an error for FREQUENCIES in general mode was corrected on March 1, 1977. Alphanumeric values were incorrectly sorted when the frequency table was generated.

On March 1, 1977 revision 3 of P-STAT 3.07 was installed. It made the following corrections:

1. Both NO.MATCH in an up/down JOIN and binary masks in Max(1-24(10)) now work.
2. In DATA, set "* MISS1-2-3" as a complicated function of the input record length is now possible.
3. In DISCRIM, incorrect function files are no longer produced.

TSO STATISTICS

These statistics are for M.I.T. usage of TSO only.

Number of LOGON sessions	6,759
Total CONNECT time	3,612 hours
Average CONNECT time per session	32 minutes
Total CPU time	16.11 hours
Average CPU time per session	0.14 minutes
Total disk I/O operations	8.3 million
Average disk I/O operations per session	1,231
Total teleprocessing I/O operations	1.9 million
Average teleprocessing I/O operations per session	285
Total K byte minutes of interactive memory use	0.62 million
Average K byte minutes per session	92

5/370 PERFORMANCE

With the problems we experienced in February, we're grateful it's the shortest month of the year. Over the first weekend we experienced a hardware failure which resulted in the 370/168 being unavailable for 25 consecutive hours. Actually, it was available for the engineers, but it took that length of time to identify and correct the problem.

Total downtime and unscheduled interruptions in February exceeded the previous three months combined. If we were due for a bad month, we're glad it came when there were only twenty eight days.

	January 1977	February 1977
Hours available (JPS and TSO)	382.50	361.43
Daily average of hours available	12.5	12.9
Hours of unattended TSO available	295.0	215.0
Daily average of hours of unattended TSO	9.0	7.5
Hours lost due to failures	14.5	46.5
Hardware	14.5	44.5
Software	.0	2.0
Number of failures	8.0	19.0
Mean time between failures	84.6 hours	30.5 hours
Mean time to repair	92.5 minutes	146.8 minutes
Actual/scheduled uptime	98.8%	93.4%

14 JOB PROCESSING SYSTEM STATISTICS

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	15,577	
Number of M.I.T. job steps processed on the 168	54,221	
Total CPU time used	67	hours
Average CPU time used per job	0.29	minutes
*Total run time used (estimated)	218	hours
Average run time per job	0.84	minutes
Total number of I/O operations by type		
CTC	5.6	million
Tape	2.7	million
Disk	15.1	million
Average number of I/O operations per job		
CTC	409	
Tape	1,321	
Disk	1,072	
Total memory minutes used	3.2	million k bytes
Average memory minutes used per job	228	K bytes
Total number of lines printed	20.5	million
Average number of lines printed per job	1,347	
Total number of cards read	4.4	million
Average number of cards read per job	280	
Total number of cards punched	0.43	million
Average number of cards punched per job	409	
Total number of Remote Job I/O units	3.0	thousand
Average number of RJ I/O units per job	335	
Total number of setup jobs	3,602	
Total number of setup tapes	7,140	
Total number of setup disks	4,669	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP system.

TURN-AROUND STATISTICS FOR FEBRUARY 1977

The following table reports turn-around figures for the Job Processing System for the month of February, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		6,959	7,591	306	15,582
Average Turn-around		42 mins.	20 mins.	12 mins.	137 mins.
Percentage of Jobs Returned	25%	7 mins.	5 mins.	4 mins.	6 mins.
	50%	23 mins.	10 mins.	9 mins.	12 mins.
	75%	45 mins.	15 mins.	13 mins.	30 mins.
	95%	149 mins.	59 mins.	30 mins.	405 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 11,976 non-setup jobs was . . . . . .101 mins.
The average turn-around for 3,606 setup jobs was . . . . . .257 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 120K	1,819	153 mins.	12 mins.
122 - 200K	6,529	97 mins.	12 mins.
202 - 300K	4,718	164 mins.	12 mins.
302 - 400K	1,652	158 mins.	15 mins.
402 - 500K	418	177 mins.	22 mins.

TSO STATISTICS

These statistics are for M.I.T. usage of TSO only.

Number of LOGON sessions	6,665	
Total CONNECT time	3,509	hours
Average CONNECT time per session	32	minutes
Total CPU time	15.2	hours
Average CPU time per session	0.14	minutes
Total disk I/O operations	7.9	million
Average disk I/O operations per session	1,181	
Total teleprocessing I/O operations	1.7	million
Average teleprocessing I/O operations per session	265	
Total K byte minutes of interactive memory use	0.56	million
Average K byte minutes per session	85	

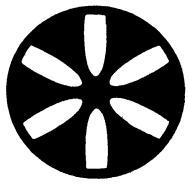

5/370 PERFORMANCE

During the first month of 1977 we had to close the facility on three occasions because of stormy conditions. However, on the positive side, we enjoyed a full month free of software related system failures. The eight unscheduled interruptions during January were all hardware related. One failure, which occurred during unattended service, accounted for ten hours of lost time. Fortunately, the other failures were identified and repaired within an hour after occurring.

	December 1976	January 1977
Hours available (JPS and TSO)	397.25	382.50
Daily average of hours available	12.8	12.5
Hours of unattended TSO available	300.0	295.0
Daily average of hours of unattended TSO	9.6	9.0
Hours lost due to failures	10.5	14.5
Hardware	10.0	14.5
Software	.5	.0
Number of failures	10.0	8.0
Mean time between failures	69.7 hours	84.6 hours
Mean time to repair	63.0 minutes	92.5 minutes
Actual/scheduled uptime	99.0%	98.8%

IL JOB PROCESSING SYSTEM STATISTICS

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	16,338	
Number of M.I.T. job steps processed on the 168	55,277	
Total CPU time used	64	hours
Average CPU time used per job	0.26	minutes
*Total run time used (estimated)	207	hours
Average run time per job	0.76	minutes
Total number of I/O operations by type		
CTC	6.3	million
Tape	2.3	million
Disk	14.3	million
Average number of I/O operations per job		
CTC	436	
Tape	1,163	
Disk	972	
Total memory minutes used	3.2	million k bytes
Average memory minutes used per job	215	K bytes
Total number of lines printed	20.6	million
Average number of lines printed per job	1,296	
Total number of cards read	4.5	million
Average number of cards read per job	276	
Total number of cards punched	0.36	million
Average number of cards punched per job	329	
Total number of Remote Job I/O units	3.6	thousand
Average number of RJ I/O units per job	360	
Total number of setup jobs	3,566	
Total number of setup tapes	7,019	
Total number of setup disks	4,568	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP system.

TURN-AROUND STATISTICS FOR JANUARY 1977

The following table reports turn-around figures for the Job Processing System for the month of January, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		6,511	8,928	261	16,338
Average Turn-around		37 mins.	19 mins.	16 mins.	115 mins.
Percentage of Jobs Returned	25%	7 mins.	5 mins.	5 mins.	6 mins.
	50%	14 mins.	9 mins.	10 mins.	11 mins.
	75%	42 mins.	14 mins.	15 mins.	28 mins.
	95%	134 mins.	59 mins.	59 mins.	195 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 12,772 non-setup jobs was . . . . . 82 mins.
The average turn-around for 3,566 setup jobs was . . . . . .232 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 110K	1,722	193 mins.	11 mins.
112 - 200K	6,560	81 mins.	11 mins.
202 - 300K	5,270	98 mins.	11 mins.
302 - 400K	1,842	147 mins.	13 mins.
402 - 500K	396	310 mins.	23 mins.

ARPANET NEWS

by Jean Bow

Jean Bow of M.I.T. Information Processing Services (User Services) is the new liaison for the ARPA Network. She has taken over this position from George Williams who has transferred to the Systems Programming Group. Jean will answer questions regarding software packages available on M.I.T.'s Multics, access to the system, prices for services and other pertinent information. If you are already a Multics user, and are not at M.I.T., you can send Multics mail to Jean (Bow IPCUS) to obtain information on how to use Multics remotely via the ARPANET. If you are a potential remote user and anxious to find out what software packages are available on Multics, please call Jean at (617) 253-7185 or contact her by U.S. mail at the following address:

Jean Bow
User Services
M.I.T. Information Processing Services
Room 39-425
77 Massachusetts Avenue
Cambridge, MA 02139

Revisions and enhancements to the ARPA telecommunications network are being contemplated. To assist in the conceptual design of the above, the Network Information Center would like to have your opinions, comments and criticisms concerning current ARPANET features and operation, and ideas for the next network design.

Please send your contributions

by Network Mail to FUTURE-NET@OFFICE-1

or

by U.S. mail to Elizabeth Feinler
Network Information Center, J2021
Stanford Research Institute
Menlo Park, California 94205

or to Jean Bow by Multics or U.S. mail.

ERROR IN "NEW PARAMETERS" ARTICLE

by Toni Calavas

In last month's article, "New Parameters for exec_com and do", we misspelled the command used to create directories. The correct way to define the abbreviation used in the example is:

```
.ab create do "create_dir &1;sa &1 sma *.OurProj.*"
```

The abbreviation is used this way:

```
create test
```

Actually, you'll probably want to use a name other than "create", which is a standard Multics command to do something else. For example, the abbreviation could be make_dir instead. We used create only for consistency with the help file.

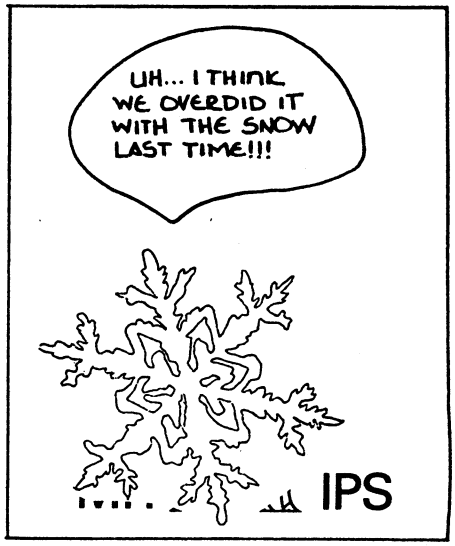

HELP FOR THE NEW YEAR

by Sue Thompson

Want to keep track of those business meetings or that dentist appointment that you keep forgetting? Type "help calendar". "calendar" prints a calendar for one month, using files that you create to list calendar events. The calendar has the month name and a two digit year at the top in big letters; each calendar box is 16 characters wide and seven characters high. The boxes will contain nothing but the number of the day in the month unless one or more files are specified in the command line.

For an example, we have made a calendar for the month of February, which includes the IPS course and seminar offered that month as well as any holidays that occur. The syntax for calendar is:

```
calendar date_string -ca- file1 file2...
```

where date_string is the first date in the month to be pictured, in this case 02/01/77, and file1 and file2, and so on, represent the items to be inserted in the calendar, holidays and IPS courses respectively. Thus, to print our February calendar, the command line looked like this:

```
calendar 02/01/77 hol.runoff cour.runoff
```

and, the result is (reduced):

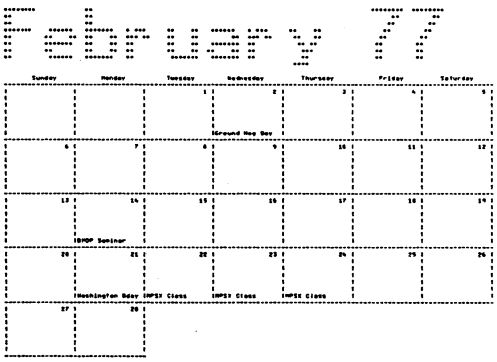

The files hol.runoff and cour.runoff looked like this:

```
date,02/02,Ground Hog Day
date,02/21,Washington Bday

date,02/14,BMDP Seminar
date,02/22,MPSX Class
date,02/23,MPSX Class
date,02/24,MPSX Class
```

Further instructions on how to define specific dates on the calendar are included in the help file, "help calendar". Have fun!

MULTICS OPERATING STATISTICS

	November	December	% Change
System hours up	580	630	+8.5
Two CPU hours	492	409	-16.9
MTBF (hours)	24:10	17:59	-25.6
CPU charged (hours)	198	198	-0.1
Memory units (million)	9.3	11.0	+19.1
Print lines (million)	11.3	12.8	+13.1
Terminal hours	12,845	13,041	+1.5
Logins	12,203	11,894	-2.5
Absentee jobs	2,280	2,219	-2.7
Disk usage (page-months)	117,111	110,421	-5.7
Projects	205	214	+4.4
Users	1,485	1,476	-0.6

MORE HINTS FOR FORTRAN USERS

Since the December installation of the MOD II version of the FORTRAN utility routines, some people have encountered the Linkage Editor message "IEW0241 WARNING-EXTERNAL SYMBOL IS DOUBLY DEFINED" when relinking an existing load module with additional or replacement routines. One procedure for such problems is to reconstruct the load module from object or source versions of the routines.

Since some load modules cannot be reconstructed easily, an alternative procedure may be preferred. The Linkage Editor can be instructed to make deletions and substitutions of the control sections within a module. The IPS staff has created a data set of REPLACE statements for all the MOD I names (IHN...), and has set read access for public use. The data set name is U.C7612.P88888.FORTMODI.REPLACE. These statements should precede an INCLUDE statement for the existing module.

The following job creates a new load module from an existing load module, replacing MOD I routines with MOD II equivalents and including new or modified FORTRAN subroutines:

```
   Job parameter statements  (see OS-9)
   ...
//RELINK EXEC FORCL
//C.SYSIN  DD *,DCB=BLKSIZE=2000
   ...
   New FORTRAN routines in source form
   ...
//L.SYSLIN  DD
//    DD  DSNAME=U.C7612.P88888.FORTMODI.REPLACE,DISP=SHR
//    DD  *
   INCLUDE   OLD(OLDMEM)
   ENTRY MAIN
//L.SYSLMOD  DD  DSNAME=U.M12345.P67890.MYLIB.LOAD(NEWMEM),DISP=OLD
//L.OLD    DD  DSNAME=U.M12345.P67890.MYLIB.LOAD,DISP=SHR
```

Note that DDname SYSLIN is required to concatenate the on-line control statements with those which are provided after the *. The first reference to SYSLIN is left blank because output from the compiler is passed to a temporary dataset associated with SYSLIN.

The following TSO commands could be used to create a revised load module with only MOD I references replaced by MOD II routines:

```
alloc file(old) dataset(mylib.load)
READY
link ('u.c7612.p88888.fortmodi.replace' *) -
      fortlib map load(mylib.load(newmem)) print(*)
ENTER CONTROL STATEMENTS -
      include old(oldmem)

END OF CONTROL STATEMENTS
READY
```

You can check the Linkage Editor MAP to be sure that all MOD I references have been replaced: MOD I names begin with IHN while IHO identifies MOD II names.

If you have any questions about this process, please contact the Programming Assistance and Information Office, Room 39-219, telephone (617) 253-4114.

Another change you may have noticed since the installation of the MOD II library is error message prefixes. The MOD II error message prefix is IHO; whereas the MOD I prefix for the same messages was IHN. The IHO error messages are listed in "IBM OS FORTRAN IV (H Extended) Compiler and Library (Mod II) Messages", order number SC28-6865-1. However, since only the error prefix has changed, the IHN codes can be used instead. For example, message IHN240, which was associated with the FORTRAN MOD I library, cites the same error as the new code IHO240. These codes can be found in "IBM OS (TSO) Terminal User's Supplement for FORTRAN IV (G1) Processor and TSO FORTRAN Prompter", SC28-6855.

PRIVATE DATASETS REVISITED

by Gerald Swanson

In the December Bulletin article, "In Quest of Private Datasets", we inadverently omitted an essential part of the example: the renaming of the dataset on the 2314 disk pack.

To rename your cataloged dataset, PV.MYDATA, on 2314 volume 234000 to a name following the Center's convention, use the following job:

```
  Job parameter statements (see OS-9)                          Col. 72
  ...                                                             |
/*SETUP UNIT=2314,ID=234000,A=XYZ                                 |
//  EXEC MITUTIL,PROG=IEHPROGM,UNIT1=2314,VOL1=234000            |
//SYSIN DD *                                                     |
 UNCATLG DSNAME=PV.MYDATA                                        |
 CATLG DSNAME=PV.M12345.P67890.MYDATA,VOL=2314=234000            |
 RENAME DSNAME=PV.MYDATA,VOL=2314=234000,                        X
              NEWNAME=PV.M12345.P67890.MYDATA
              ↑
         Col. 16
```

If you have any questions, please contact Gerald Swanson, Room 39-423, telephone (617) 253-1732.

TSO STATISTICS

These statistics are for M.I.T. usage of TSO only.

Number of LOGON sessions	6,611	
Total CONNECT time	3,489	hours
Average CONNECT time per session	32	minutes
Total CPU time	17.6	hours
Average CPU time per session	0.16	minutes
Total disk I/O operations	8.0	million
Average disk I/O operations per session	1,217	
Total teleprocessing I/O operations	1.8	million
Average teleprocessing I/O operations per session	272	
Total K byte minutes of interactive memory use	0.61	million
Average K byte minutes per session	92	

S/370 PERFORMANCE

Twenty-three days in December 1976 were free of any unscheduled interruptions, and we enjoyed another month free of major problems. Of the ten crashes, seven were due to a variety of hardware failures. One failure, which occurred during unattended service, resulted in TSO being unavailable for a seven hour period.

	November	December 1976
Hours available (JPS and TSO)	388.30	397.25
Daily average of hours available	12.9	12.8
Hours of unattended TSO available	277.5	300.0
Daily average of hours of unattended TSO	9.0	9.6
Hours lost due to failures	2.3	10.5
Hardware	1.3	10.0
Software	1.0	.5
Number of failures	10.0	10.0
Mean time between failures	66.5 hours	69.7 hours
Mean time to repair	13.1 minutes	63.0 minutes
Actual/scheduled uptime	98.7%	99.0%

JOB PROCESSING SYSTEM STATISTICS

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	16,753	
Number of M.I.T. job steps processed on the 168	57,579	
Total CPU time used	65	hours
Average CPU time used per job	0.26	minutes
*Total run time used (estimated)	210	hours
Average run time per job	0.75	minutes
Total number of I/O operations by type		
CTC	8.6	million
Tape	2.1	million
Disk	14.4	million
Average number of I/O operations per job		
CTC	576	
Tape	1,216	
Disk	950	
Total memory minutes used	3.1	million k bytes
Average memory minutes used per job	203	K bytes
Total number of lines printed	24.5	million
Average number of lines printed per job	1,486	
Total number of cards read	5.0	million
Average number of cards read per job	299	
Total number of cards punched	0.36	million
Average number of cards punched per job	306	
Total number of Remote Job I/O units	7.7	thousand
Average number of RJ I/O units per job	513	
Total number of setup jobs	3,455	
Total number of setup tapes	6,816	
Total number of setup disks	5,023	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP system.

TURN-AROUND STATISTICS FOR DECEMBER 1976

The following table reports turn-around figures for the Job Processing System for the month of December, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		6,754	9,000	374	16,754
Average Turn-around		37 mins.	18 mins.	11 mins.	104 mins.
Percentage of Jobs Returned	25%	7 mins.	5 mins.	5 mins.	6 mins.
	50%	14 mins.	10 mins.	10 mins.	11 mins.
	75%	43 mins.	14 mins.	14 mins.	28 mins.
	95%	134 mins.	59 mins.	44 mins.	195 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 13,299 non-setup jobs was . . . . . 83 mins.
The average turn-around for 3,455 setup jobs was . . . . . .183 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 100K	1,973	81 mins.	10 mins.
112 - 200K	7,772	83 mins.	11 mins.
202 - 300K	4,204	120 mins.	11 mins.
302 - 400K	1,712	132 mins.	13 mins.
402 - 500K	498	204 mins.	23 mins.

NEW PARAMETERS FOR EXEC_COM AND DO
By:Toni Calavas

The do command and active function and the exec_com command recognize two new parameters for argument substitution. The parameter &fi (where i is an integer between one and nine inclusive) is replaced by the ith through last actual arguments supplied. The parameter &n is replaced by the number of arguments supplied.

For example,

```
do "ioa_ (&f3)" a b c d e
```

prints the strings "c", "d" and "e". The command

```
do "ioa_ &f&n" a b c d e
```

prints the string "e".

The do command can also be very useful when used in conjunction with the abbrev command (which permits abbreviation of a command line), because it permits substitution of an argument. For example, the abbreviation "create" enables the user to create any directory and set access on it for all users in the project. Here's what the abbreviation looks like:

```
create do "create directory &1; sa &1 sma *.OurProj.*"
```

This is how you'd use the abbreviation:

```
create test
```

An exec_com is useful when you have a task that requires the same sequence of Multics commands, with perhaps minor variations (such as changing the segment name), that must be issued over and over.

BETTER HELP NOW
By: Toni Calavas

Have you ever wanted to print a help file stored in another directory? Well, in the past you could by using the -pathname (-pn) option, but now it is easier. We have changed the help command to accept pathnames as well as entry names.

This is how it works. If the name specified in the help command line contains a ">" character, it is interpreted as a pathname. For example, if a colleague of yours has an information segment in her (info) directory, you can print it by typing

```
help >udd>ProjectA>PersonB>info>name
```

If you have changed to her working directory, you'd print the helpfile by typing

```
help -pn name
```

Help is the command you can use to print system segments containing information about commands, subroutines, general information, system changes, etc. You can obtain more information about the kinds of help files available by typing

```
help info_segs
```

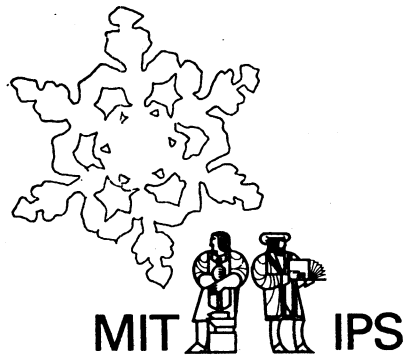

MULTICS OPERATING STATISTICS

	October	November	% Change
System hours up	637	580	-9.0
Two CPU hours	632	492	-22.2
MTBF (hours)	79:40	24:10	-69.7
CPU charged (hours)	182	198	+8.9
Memory units (million)	9.3	9.3	-1.0
Print lines (million)	10.6	11.3	+6.9
Terminal hours	13,509	12,845	-4.9
Logins	12,554	12,203	-2.8
Absentee jobs	2,009	2,280	+13.5
Disk usage (page-months)	106,749	117,111	+9.7
Projects	204	205	+0.5
Users	1,474	1,485	+0.7

OS 370

IMSL EDITION 5 INSTALLED

The IMSL subroutine library has been updated recently to Edition 5. Highlights include:

-subroutines for the generalized eigenvector problem

-improved cubic spline routines

-an improved linear equation solver for indefinite symmetric matrices

-a new nonlinear least-squares routine

-a new routine to find the minimum of a function of n variables

-ARIMA time-series model analysis

-complete factor analysis including estimation and rotation options

-regression model selection via "leaps and bounds" algorithm

This change removed thirty-two routines, added fifty-six others, and modified fifty-eight. Users should be aware that a more efficient random number generator, GGUB, is now used in place of GGUl, requiring a change of parameters in fourteen routines which called GGUl. Also, the dimension of the parameter TEMP for the routine BECORO has been modified. A copy of the IMSL update newsletter is posted in the Programming Assistance and Information Office, 39-219. Except for the routines listed there, all other aspects of the library use remain as described in AP-9.

FORTRAN UTILITIES UPGRADED

by Art Anger

The FORTRAN library of utility functions, SYS1.FORTLIB, was scheduled for upgrading to the Mod-II version on 20 December. This change corrects a number of potential errors in special input or output situations and may offer a few improvements in efficiency. Users of Version 6 of SPSS have been exercising this library for many months, with no difficulties attributable to it.

Since there are no changes in intended capability, most users will not need to make any adaptations for this library. The two situations where it will make a difference lie in the use of previously linked modules. A completely linked module will continue to work, except that error reports may be greatly abbreviated once the Mod I dynamic error handler is removed. A partially linked module (such as the dynamically loaded routines of EISPACK) can encounter difficulties, because inter-utility reference names do change with the version. For instance, SQRT was an alias for IHNSQRT, but now points to IHOSQRT. A user program referring to SQRT will always get the current (IHO) version, but an old module containing a utility routine which also uses the square-root routine will still refer to the old (IHN) version. The cure for both types of difficulty is to recreate any load modules containing IHN (or even older IHC) control sections. (An alternative is to replace them by means of appropriate Linkage Editor commands.)

For assistance with special cases, bring your load-module documentation to the Programming Assistance and Information Office, Room 39-219 (Ext. 3-4114).

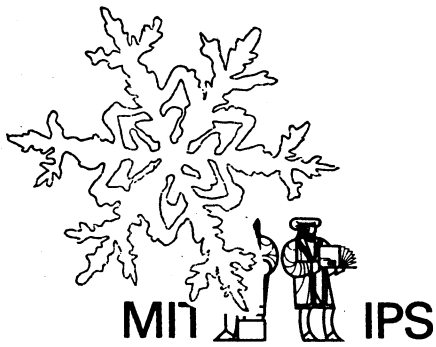

In Quest of Private Datasets

Private 2314 cataloged datasets (those beginning with PV - for example, PV.M12345.P67890) are now listed in USERMAP. USERMAP is a listing of all on-line dataset names and information about them; USERMAP will now also contain a listing of the off-line cataloged dataset names (datasets on private volumes). A listing of information printed in the USERMAP can be found in Center memo OS-2, which is available in the Publications Office, Room 39-484; additional information can be found by typing "help usermap".

The naming convention for PV datasets is:

PV.projno.progno.furtherid

Note that Center policy requires that PV catalog entries conform to this naming convention. We may delete names not conforming to the convention at any time.

To rename your cataloged dataset, PV.MYDATA, on 2314 volume 234000 to a name which follows our convention, use the following job:

```
    Job parameter statements (see OS-9)
    ...
//    EXEC MITUTIL,PROG=IEHPROGM
//SYSIN  DD  *
   UNCATLG DSNAME=PV.MYDATA
   CATLG   DSNAME=PV.M12345.P67890.MYDATA,VOL=2314=234000
/*
```

COMPUTATION TIMING
by Art Anger

CAUTION: Although it is safe to say that the IBM 370/168 executes 3 million instructions per second, trying to determine "the execution time" of any particular sequence of operations can be dangerous to your mental health. Reasons follow.

A multi-tasking operating system can interrupt any but the most sensitive (system) tasks for any of a great many reasons, only one of which is because the executing program has "completed its turn" in the rotating queue of active work. Tests with a microsecond timing routine show, in fact, that uninterrupted intervals as short as 50 microseconds (about 75 instructions) are hard to find on a busy afternoon. This conclusion is based on reports of time elapsed during each of many repetitions of the same sequence of steps; the minimum duration can occur as seldom as once in ten or twenty trials.

Another measurement, the time billed to the user program, is much more consistent, but still fluctuates as much as 10 or 20 percent with other activity on the system. Some system actions are performed for the direct benefit of the job whose primary computation is interrupted, but they are not synchronized with the instructions being timed. In addition, a small portion of other interruption periods is unavoidably charged to the job in progress. The reporting routine TIMING attempts to mask this variability by reporting only hundredths of seconds, but thereby introduces additional variation through the rounding process.

A third frustration is that the speed-controlling factor for most computations is storage access, not its operation complexity. On the 370/168, main storage is read in 8-byte groups into the high-speed "cache" memory accessible to the execution logic. Execution of a long sequence of instructions stored in successive locations can proceed only at the average speed of main-memory fetching, which can be delayed briefly by input or output accesses elsewhere in the same storage bank; duration will be roughly proportional to the number of double words of instructions plus the number of data-element accesses. Execution of a short loop is governed by its ability or inability to fit completely into cache memory, which can be influenced by its location relative to double-word boundaries and by the allocation algorithm for cache storage.

For statistically trained programmers who wish to investigate these claims further, memo AP-37 describes TIMING, for billed time in centiseconds, and INTVL, for elapsed time in microseconds. For others who are merely interested in having their programs run as cheaply as possible, memo PP-14 describes several ways to revise a program to reduce execution time.

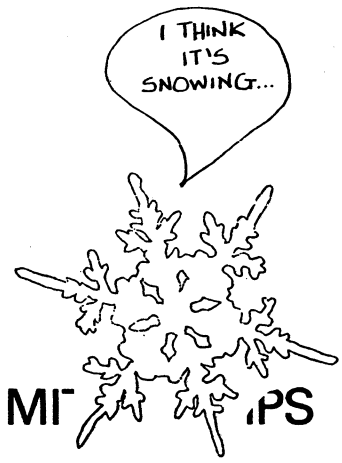

BACK UP YOUR 2314 DISK PACKS WITH FAST DUMP RESTORE

The Fast Dump Restore procedure (FDR) can be used to back up your private 2314 disk packs to tape and to restore them from tape to disk. FDR is more efficient than the IBM utility IEHDASDR.

For instance, to dump the 2314 disk pack 234000 to standard-labeled tape 001234, you would use the following job stream:

```
   Job parameter statements (see OS-9)
   . . .
/*SETUP UNIT=2314,ID=234000,A=XYZ
/*SETUP UNIT=TAPE9,ID=(001234,RING,SAVE,SL),A=ABC
//      EXEC FDR,DISK=234000,TAPE=001234
/*
```

To restore the dumped disk pack from tape 001234 to a new 2314 disk pack, 234111, you would use this job stream:

```
   Job parameter statements (see OS-9)
   . . .
/*SETUP UNIT=2314,ID=234111,A=DEF
/*SETUP UNIT=TAPE9,ID=(001234,NORING,SAVE,SL),A=ABC
//    EXEC FDR,PRM=N,DISK=234111,TAPE=001234
//TAPE1 DD DSN=FDR.$234000
/*
```

FDR also has special functions which allow you to dump and restore specific datasets. For more information, contact Gerald Swanson, Room 423 (telephone 253-1732).

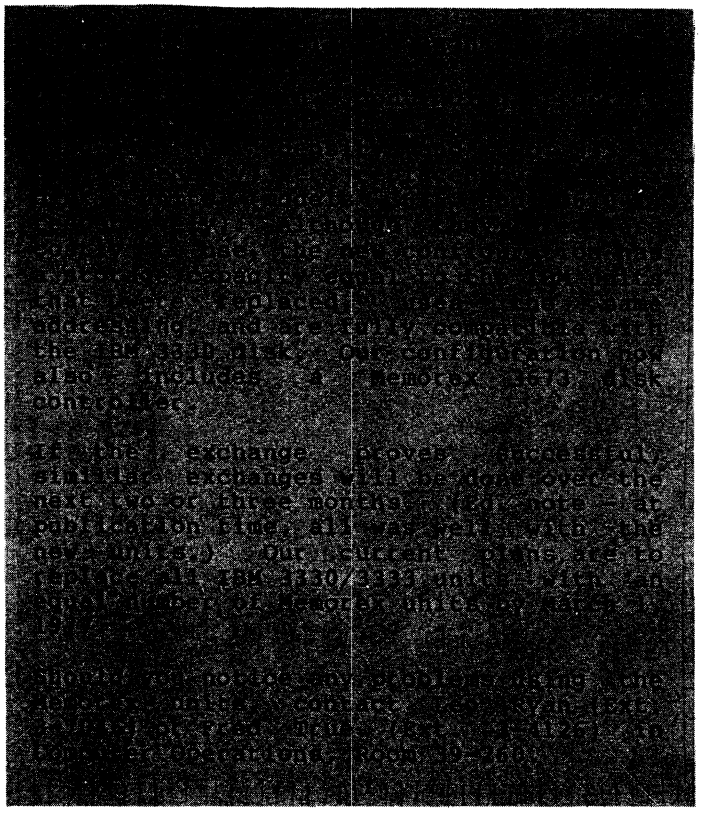

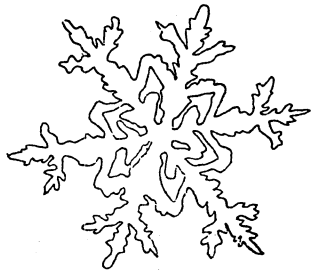

TSO STATISTICS

These statistics are for M.I.T. usage of TSO only.

Number of LOGON sessions	8,028	
Total CONNECT time	4,494	hours
Average CONNECT time per session	34	minutes
Total CPU time	21.1	hours
Average CPU time per session	0.16	minutes
Total disk I/O operations	9.7	million
Average disk I/O operations per session	1,210	
Total teleprocessing I/O operations	2.2	million
Average teleprocessing I/O operations per session	281	
Total K byte minutes of interactive memory use	0.72	million
Average K byte minutes per session	90	

S/370 PERFORMANCE

Ten unscheduled service interruptions occurred during November, equally distributed between software and hardware.

We also had one planned ten-hour shutdown to allow time for an exchange of a portion of our 3330 disk drives.

	October	November 1976
Hours available (JPS and TSO)	402.50	388.30
Daily average of hours available	12.9	12.9
Hours of unattended TSO available	305.4	277.5
Daily average of hours of unattended TSO	9.0	9.0
Hours lost due to failures	4.6	2.3
Hardware	3.5	1.3
Software	1.1	1.0
Number of failures	12.0	10.0
Mean time between failures	58.0 hours	66.5 hours
Mean time to repair	23.0 minutes	13.1 minutes
Actual/scheduled uptime	99.1%	98.7%

JOB PROCESSING SYSTEM STATISTICS

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	19,173	
Number of M.I.T. job steps processed on the 168	63,046	
Total CPU time used	84	hours
Average CPU time used per job	0.29	minutes
*Total run time used (estimated)	261	hours
Average run time per job	0.82	minutes
Total number of I/O operations by type		
CTC	7.8	million
Tape	2.6	million
Disk	18.0	million
Average number of I/O operations per job		
CTC	459	
Tape	1,380	
Disk	1,036	
Total memory minutes used	3.7	million k bytes
Average memory minutes used per job	216	K bytes
Total number of lines printed	24.1	million
Average number of lines printed per job	1,281	
Total number of cards read	5.3	million
Average number of cards read per job	278	
Total number of cards punched	0.39	million
Average number of cards punched per job	304	
Total number of Remote Job I/O units	5.9	thousand
Average number of RJ I/O units per job	423	
Total number of setup jobs	3,732	
Total number of setup tapes	7,310	
Total number of setup disks	5,549	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP system.

TURN-AROUND STATISTICS FOR NOVEMBER 1976

The following table reports turn-around figures for the Job Processing System for the month of November, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		8,687	9,128	471	19,175
Average Turn-around		45 mins.	17 mins.	27 mins.	148 mins.
Percentage of Jobs Returned	25%	7 mins.	5 mins.	4 mins.	6 mins.
	50%	14 mins.	9 mins.	9 mins.	12 mins.
	75%	56 mins.	14 mins.	14 mins.	29 mins.
	95%	164 mins.	59 mins.	44 mins.	525 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 15,443 non-setup jobs was . . . . . .117 mins.
The average turn-around for 3,732 setup jobs was . . . . . .274 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 110K	2,080	149 mins.	11 mins.
112 - 200K	8,748	124 mins.	11 mins.
202 - 300K	5,118	144 mins.	12 mins.
302 - 400K	2,257	166 mins.	13 mins.
402 - 500K	485	373 mins.	25 mins.

S/370 PERFORMANCE IN OCTOBER 1976

We were able to have TSO available for over 22.5 hours on twenty-six days this month. While we feel that is a respectable average, we would still like one month free of any unscheduled interruptions. In October, there were twelve; six of each for the hardware and software.

	September	October 1976
Hours available (JPS and TSO)	390.95	402.50
Daily average of hours available	13.03	12.9
Hours of unattended TSO available	275.3	305.4
Daily average of hours of unattended TSO	9.1	9.0
Hours lost due to failures	4.5	4.6
Hardware	3.2	3.5
Software	1.3	1.1
Number of failures	17.0	12.0
Mean time between failures	37.0 hours	58.0 hours
Mean time to repair	29.2 minutes	23.0 minutes
Actual/scheduled uptime	98.7%	99.1%

TSO STATISTICS FOR OCTOBER 1976

These statistics are for M.I.T. usage of TSO only.

Number of LOGON sessions	6,888	
Total CONNECT time	3,942	hours
Average CONNECT time per session	34	minutes
Total CPU time	15.9	hours
Average CPU time per session	0.14	minutes
Total disk I/O operations	7.3	million
Average disk I/O operations per session	1,053	
Total teleprocessing I/O operations	1.9	million
Average teleprocessing I/O operations per session	282	
Total K byte minutes of interactive memory use	0.55	million
Average K byte minutes per session	80	

JOB PROCESSING SYSTEM STATISTICS FOR OCTOBER

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	14,680	
Number of M.I.T. job steps processed on the 168	55,683	
Total CPU time used	63	hours
Average CPU time used per job	0.29	minutes
*Total run time used (estimated)	180	hours
Average run time per job	0.73	minutes
Total number of I/O operations by type		
CTC	5.6	million
Tape	1.8	million
Disk	11.7	million
Average number of I/O operations per job		
CTC	427	
Tape	1,055	
Disk	878	
Total memory minutes used	2.5	million k bytes
Average memory minutes used per job	184	K bytes
Total number of lines printed	18.4	million
Average number of lines printed per job	1,282	
Total number of cards read	4.3	million
Average number of cards read per job	292	
Total number of cards punched	0.42	million
Average number of cards punched per job	484	
Total number of Remote Job I/O units	5.2	thousand
Average number of RJ I/O units per job	648	
Total number of setup jobs	3,169	
Total number of setup tapes	7,214	
Total number of setup disks	5,014	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP system.

TURN AROUND STATISTICS FOR OCTOBER

The following table reports turn-around figures for the Job Processing System for the month of October, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		7,307	6,503	170	14,680
Average Turn-around		28 mins.	11 mins.	10 mins.	133 mins.
Percentage of Jobs Returned	25%	6 mins.	4 mins.	4 mins.	5 mins.
	50%	12 mins.	9 mins.	9 mins.	11 mins.
	75%	28 mins.	13 mins.	14 mins.	27 mins.
	95%	104 mins.	30 mins.	44 mins.	420 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 11,511 non-setup jobs was .113 mins.
The average turn-around for 3,169 setup jobs was .203 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 100K	1,768	171 mins.	10 mins.
102 - 200K	6,651	112 mins.	10 mins.
202 - 300K	3,964	119 mins.	11 mins.
302 - 400K	1,475	154 mins.	11 mins.
402 - 500K	375	286 mins.	14 mins.

IMSL Documentation Error

The IMSL manual, Computer Subroutine Libraries in Mathematics and Statistics (Library 1, Edition 5) contains an error in the description of subroutine OFPRIN, which performs principal components analysis. It incorrectly states the dimensions of the output, A, the matrix of factor loadings. Although on output, the matrix will be NV by NF, this space is used by EIGRS to compute all eigenvectors, so that A must be dimensioned NV by NV.

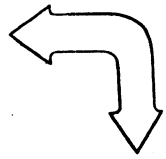

S/370 Performance in September, 1976

September was, fortunately, a fairly uneventful month. We did have problems installing a new disk drive, (six crashes), but that finally was resolved. Of the remaining crashes, seven were hardware (four on one day for one device) and four were software. We always have hopes the hardware crashes will go away, but a software crash every week seems to be unavoidable, at least on this operating system.

	August	September 1976
Hours available (JPS and TSO)	452.25	390.95
Daily average of hours available	14.6	13.03
Hours of unattended TSO available	252.4	275.5
Daily average of hours of unattended TSO	8.5	9.1
Hours lost due to failures	9.3	4.5
Hardware	6.2	3.2
Software	2.1	1.3
Number of failures	17.0	17.0
Mean time between failures	41.5 hours	37.0 hours
Mean time to repair	32.1 minutes	29.2 minutes
Actual/scheduled uptime	98.9%	98.7%

TSO Statistics for September, 1976

These statistics are for M.I.T. usage of TSO only.

Number of LOGON sessions	6,056	
Total CONNECT time	3,462	hours
Average CONNECT time per session	34	minutes
Total CPU time	12.9	hours
Average CPU time per session	0.13	minutes
Total disk I/O operations	6.2	million
Average disk I/O operations per session	1,027	
Total teleprocessing I/O operations	1.6	million
Average teleprocessing I/O operations per session	269	
Total K byte minutes of interactive memory use	0.46	million
Average K byte minutes per session	75	

JOB PROCESSING SYSTEM STATISTICS FOR SEPTEMBER, 1976

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	13,033	
Number of M.I.T. job steps processed on the 168	55,467	
Total CPU time used	55	hours
Average CPU time used per job	0.29	minutes
*Total run time used (estimated)	214	hours
Average run time per job	0.98	minutes
Total number of I/O operations by type		
CTC	6.6	million
Tape	2.5	million
Disk	16.0	million
Average number of I/O operations per job		
CTC	576	
Tape	1,455	
Disk	1,358	
Total memory minutes used	2.8	million k bytes
Average memory minutes used per job	152	K bytes
Total number of lines printed	20.4	million
Average number of lines printed per job	1,607	
Total number of cards read	3.6	million
Average number of cards read per job	274	
Total number of cards punched	0.38	million
Average number of cards punched per job	463	
Total number of Remote Job I/O units	1.9	thousand
Average number of RJ I/O units per job	76	
Total number of setup jobs	3,257	
Total number of setup tapes	7,989	
Total number of setup disks	4,859	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP system.

TURN-AROUND STATISTICS FOR SEPTEMBER 1976

The following table reports turn-around figures for the Job Processing System for the month of September, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		6,143	6,064	198	13,026
Average Turn-around		39 mins.	14 mins.	13 mins.	145 mins.
Percentage of Jobs Returned	25%	6 mins.	5 mins.	5 mins.	6 mins.
	50%	13 mins.	9 mins.	10 mins.	11 mins.
	75%	41 mins.	14 mins.	24 mins.	29 mins.
	95%	179 mins.	45 mins.	45 mins.	450 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 9,771 non-setup jobs was . . . . . 127 mins.
The average turn-around for 3,255 setup jobs was . . . . . 197 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 100K	1,865	185 mins.	11 mins.
102 - 200K	5,713	120 mins.	11 mins.
202 - 300K	3,073	139 mins.	11 mins.
302 - 400K	1,673	157 mins.	13 mins.
402 - 500K	329	264 mins.	15 mins.

Commonwealth Scientific and Industrial Research Organization, Australia

Annual Report, Division of Computing Research. 1974-75.

Honeywell Information Systems, Incorporated

Multics PL/I Reference Manual. AM83, Revision 0. 1976.

IBM Research Center

On Stability and Accuracy of Numerical Integration Methods for Stiff Differential Equations. 1976.

On the Evaluation of Queries on a Relational Data Base System. 1976.

Modular Multiplexing Computer Interface System. 1976.

Concurrency of Operations on B-Trees. 1976.

A Comparison of Four Methods for the Evaluation of Queries in a Relational Data Base System. 1976.

Relational Data Base Management Systems. 1976.

Indiana University

The Relevance of Relevance. 1976.

Mathematisch Centrum

Exercises in Denotational Semantics. 1976.

On Integer Arithmetic Progressions of Length Four. 1976.

On a Conjecture of Erdos Concerning Sums of Powers of Integers. 1976.

Stabilized Runge-Kutta Methods for Second Order Differential Equations Without First Derivatives. 1975.

The Solution of the Order Equations of a Four-Point, Fourth Order, Two-Step Runge-Kutta Method. 1976.

Messerschmitt-Bolkow-Blohm

Digital Terrain Model - MBB-ICES DGM, An ICES Subsystem for Storing, Evaluation, and Graphical Representation of Digital Terrain Mode. (Descriptive brochure) 1976.

M.I.T. - Laboratory for Computer Science

A Secure and Flexible Model of Process Initiation for a Computer Utility. 1976.

Processor Multiplexing in a Layered Operating System. 1976.

National Physical Laboratory

A Simple Network Interacting Programs' Executive (SNIPE). 1976.

Petrocelli/Charter

Top-Down Structured Programming Techniques. 1975.

Society for Industrial and Applied Mathematics

Bayesian Statistics, A Review. 1972.

Methods for Solving Systems of Nonlinear Equations. 1974.

Numerical Solution of Two Point Boundary Value Problems. 1976.

Stanford University

The State of the Art of Computer Programming. 1976.

U.S. Department of Commerce

Aids for COBOL Program Conversion (FIPS Publication 21 to FIPS Publication 21-1) FIPS Publication 43. 1975.

University of California, Berkeley

The Cartographic Data Base in SEEDIS. 1975.

University of California, Livermore

GEARV: A Vectorized Ordinary Differential Equation Solver. 1975.

University of Maryland

Further Programs for the Solution of Large Sparse Systems of Linear Equations. 1975.

Design of an Interactive Direct-Execution System. 1975.

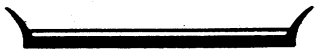

MULTICS OPERATING SYSTEM STATISTICS FOR AUGUST 1976

	July	August	% Change
System hours up	616	652	+6.0
Two CPU hours	570	628	+10.3
MTBF (hours)	30:47	59:18	+92.6
CPU charged (hours)	194	221	+13.9
Memory units (million)	8.7	10.0	+14.0
Print lines (million)	10.5	12.4	+18.4
Terminal hours	12,175	13,566	+11.4
Logins	9,914	11,201	+13.0
Absentee jobs	2,250	2,927	+30.1
Disk usage (page-months)	100,231	114,593	+14.3
Projects	197	198	+0.5
Users	1,418	1,301	-8.3

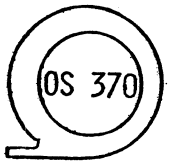

PV Catalog Entries -- A Warning

In August we ran into trouble because the private volume (PV) catalog was full. On investigating we found a number of entries that we knew could be uncataloged and were able to temporarily alleviate the problem.

We also found that users were not following our naming conventions that the second and third qualification levels must be project and programmer numbers. Please understand that in order to manage this catalog we must have that information. Accordingly any catalog entry which violates this convention will be "fair game" (for deletion) the next time the catalog fills up.

★★★★★★

TCAM Translation Table Changes

If you are a TSO user with an ASCII terminal, you might be interested in the recent changes to the TCAM (Telecommunication Access Method) translate tables. During September we installed corrections to the TCAM translation table for ASCII terminals so that:

1) It is consistent with other uses of ASCII (especially Multics);

2) It is consistent with the character set on the off-line printer's PT train; and

3) It implements more of the character set common to both EBCDIC and ASCII.

Especially important to PL/I users is the change from a tilde (~) to a caret (^) for entering a not sign (¬) on ASCII terminals.

The changes are all fully described in a recent Center memo TOS-T which is available in the Publications Office (Room 39-484). The memo also explains how to use the FAC (Format Alternate Characters) program to convert datasets containing brackets, braces, and backslashes and how to re-link-edit Tektronics load modules.

If you have any problems bring them to the PAI Office (Room 39-219).

★★
★★

S/370 Performance in August 1976

This month we averted any damage hurricane Belle might have done by shuting down the system for twelve hours. The facility withstood the wind and rain, and all equipment powered up without problem the next morning.

Twenty days were free of unscheduled service interruptions; seventeen failures, nine hardware and eight software, occurred over an eleven day period.

	July 1976	August 1976
Hours available (JPS and TSO)	449.07	452.25
Daily average of hours available	14.5	14.6
Hours of unattended TSO available	279.5	252.4
Daily average of hours of unattended TSO	9.0	8.5
Hours lost due to failures	4.7	9.3
Hardware	3.5	6.2
Software	1.2	2.1
Number of failures	10.0	17.0
Mean time between failures	72.8 hours	41.5 hours
Mean time to repair	28.6 minutes	32.1 minutes
Actual/scheduled uptime	99%	98.9%

TSO Statistics for August 1976

These statistics are for M.I.T. usage of TSO only.

Number of LOGON sessions	6,265
Total CONNECT time	3,548 hours
Average CONNECT time per session	34 minutes
Total CPU time	15.5 hours
Average CPU time per session	0.15 minutes
Total disk I/O operations	7.1 million
Average disk I/O operations per session	1,132
Total teleprocessing I/O operations	1.6 million
Average teleprocessing I/O operations per session	261
Total K byte minutes of interactive memory use	0.52 million
Average K byte minutes per session	83

Job Processing System Statistics for August 1976

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	18,163	
Number of M.I.T. job steps processed on the 168	68,069	
Total CPU time used	103	hours
Average CPU time used per job	0.38	minutes
*Total run time used (estimated)	244	hours
Average run time per job	0.81	minutes
Total number of I/O operations by type		
CTC	7.6	million
Tape	3.4	million
Disk	13.1	million
Average number of I/O operations per job		
CTC	471	
Tape	1,433	
Disk	798	
Total memory minutes used	3.4	million K bytes
Average memory minutes used per job	206	K bytes
Total number of lines printed	24.1	million
Average number of lines printed per job	1,357	
Total number of cards read	5.4	million
Average number of cards read per job	300	
Total number of cards punched	0.58	million
Average number of cards punched per job	441	
Total number of Remote Job I/O units	1.6	thousand
Average number of RJ I/O units per job	93	
Total number of setup jobs	4,179	
Total number of setup tapes	3,486	
Total number of setup disks	2,242	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP or MFT system.

TURN-AROUND STATISTICS FOR AUGUST 1976

The following table reports turn-around figures for the Job Processing System for the month of August, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		7,825	9,092	317	18,164
Average Turn-around		41 mins.	15 mins.	13 mins.	145 mins.
Percentage of Jobs Returned	25%	7 mins.	5 mins.	5 mins.	6 mins.
	50%	14 mins.	10 mins.	9 mins.	12 mins.
	75%	44 mins.	14 mins.	14 mins.	29 mins.
	95%	179 mins.	59 mins.	44 mins.	615 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 13,985 non-setup jobs was . . . . . .130 mins.
The average turn-around for 4,179 setup jobs was . . . . . .193 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 100K	1,796	136 mins.	11 mins.
102 - 200K	8,361	107 mins.	11 mins.
202 - 300K	4,780	157 mins.	12 mins.
302 - 400K	1,919	188 mins.	13 mins.
402 - 500K	883	291 mins.	14 mins.

Relief for Sore Arms

For those of our users who've been moaning under the weight of the MPM for years, (partial) relief is here! It takes the form of the Multics Pocket Guide, Commands and Active Functions (HIS order number AW17). Believe it or not, this tiny (7 1/2" x 3 3/4" x 1/4") manual summarizes all the Multics commands and active functions available, shows correct usage, and lists control arguments and optional arguments which may be used. In addition, it provides a brief definition of some Multics terminology and descriptions of the formats used. Perhaps the most valuable sections are those that list commands and active functions according to the type of functions they perform.

Some words of warning: the manual is intended for persons familiar with Multics command conventions; it is not (and wasn't intended to be) a primer or tutorial. If you aren't familiar with Multics, we suggest you first read the Multics Users' Guide (AL40), and perhaps some of the Center's MS series of memos.

One of our users has made this book even more useful for his staff by having them tape the Center's orange "Telephone Numbers to Reach IPC Time Sharing Services" cards (GI-26) inside the front cover of each book. We think that's a great idea.

If you want copies of this quick reference, you can buy them for $2.10 at the Publications Office, Room 39-484 (Telephone 253-6325). The other references mentioned above may also be obtained there.

Fine Tuning in the Multics Charging System

An important Multics feature is the system's ability to serve a wide spectrum of small and large users. It provides mechanisms to allow a finer degree of usage measurement and greater equity in costs as they relate to actual resources consumed. For instance, separate charges are made for processor (CPU) and memory.

CPU

The H6180 CPU usage excludes most system overhead functions such as processing interrupts and page faults. CPU charges range from $75 to $225 per hour depending on shift and whether usage is interactive or absentee.

Memory

The number of memory units charged to a process is an approximation of the number of page faults (references to secondary memory) that a program would take in a real memory consisting of 1000 pages (each page consists of 1024 36-bit (4-character) words). This figure is affected somewhat by system load, but our measurements show that the variation is less than the variation in the number of actual page faults. The number of memory units charged to a mixture of various commands can range from one per four page faults to one per thirty page faults. However, the number of memory units charged to an "average" program under medium system load is about one per twenty page faults. Memory charges range from $4 to $11 per 1000 memory units (Kmu) depending on shift and whether usage is interactive or absentee.

Example

Your console's ready message consists of the letter "r" followed by the time of day in hours and minutes (HHMM) and three numbers that reflect system resource usage. If the ready message reads

```
r 1741 2.790 96.238 1536
```

your program has spent 2.790 seconds of CPU time, 96.238 memory units, and 1536 page faults. Since 1741 (5:41 p.m.) is considered first shift, a quick estimate using the rate table listed in GI-7 shows the total cost is $1.23, which can be verified by using the command "estimate_bill".

Note that page faults are not used for charging.

MULTICS OPERATING STATISTICS FOR JUNE 1976

	May	June	% Change
System hours up	534	707	+32.3
Two CPU hours	388	633	+63.2
MTBF (hours)	15:42	19:38	+25.0
CPU charged (hours)	269	286	+6.4
Memory units (million)	12.0	13.2	+10.8
Print lines (million)	9.8	12.0	+22.0
Terminal hours	15,134	15,815	+4.5
Logins	13,921	12,853	-7.7
Absentee jobs	2,659	3,197	+20.2
Disk usage (page-months)	123,826	161,407	+30.4
Projects	220	223	+1.4
Users	1,959	1,543	-21.2

ALL AND EVERYTHING ABOUT DISK DATASETS

Having trouble with your disk datasets? Are you a new user of our 370/168? Then you should acquire a copy of the new memo, "Disk Datasets," OS-2. This document contains almost everything you need to know about allocating, using, and managing datasets on both the on-line disk packs and private 2314 storage. The memo is also a central source of information for the occasional user of utility programs. Simple TSO commands which perform the same task are included where applicable.

As a reminder, all users should periodically back up their on-line datasets to tape since the Center does not guarantee backup. The following job stream will back up all cataloged on-line datasets with the project and programmer numbers M12345.P67890 to standard-labeled tape number 999999:

```
        Job Parameter Statements
        . . .
/*SETUP UNIT=TAPE9,ID=(999999,RING,SAVE,SL),A=ZZZ
//      EXEC MITUTIL,PROG=IEHMOVE,UNIT1=TAPE9,VOL1=999999
//SYSIN DD *
 COPY DSGROUP=U.M12345.P67890,TO=TAPE9=999999
/*
```

For further information, see Center memos OS-1 and OS-2.

CONTOURING ON CALCOMP

A CalComp version of the contouring subroutine FNCON1 has been placed in the CalComp library SYS5.PLOT.SUBR. A description of the subroutine can be found in Appendix F of AP-36. There are, however, two changes from the latter description that are necessary.

1) The arrays X and Y must have dimensions of at least NXA+2 and NYA+2, respectively.

2) Scale factors must be provided in the calling program, either by appropriate calls to SCALE or by explicit assignments of the forms:

```
X(NXA+1) = FIRSTU       Y(NYA+1) = FIRSTV
X(NXA+2) = DELTAU       Y(NYA+2) = DELTAV
```

where NXA and NYA are the number of grid divisions actually being used.

(continued on page 9)

(FNCON1 continued)

You may submit a batch job of the following format to execute FNCON1:

```
        Job parameter statement (see OS-9)
        . . .
//      EXEC    FORCGO,LIBRARY='SYS5.PLOT.SUBR'
//C.SYSIN  DD    *,DCB=BLKSIZE=2000
        . . .
        Program
        . . .
//G.SYSIN  DD    *,DCB=BLKSIZE=2000
        . . .
        Data
        . . .
/*
```

To use FNCON1 from TSO with the CalComp previewing routines, specify the library in the LIB parameter of the LINK or LOADGO command following the Tektronix libraries:

```
LOADGO PROG FORTLIB LIB('SYS5.TLIB4010.LOAD','SYS5.TLIBPLOT.LOAD','SYS5.PLOT.SUBR')
```

Please report any problems to Mark Froimowitz, Room 39-465, Telephone 253-7729.

S/370 Performance in June 1976

Performance this month was much improved over a very bad month in May, but there were still eleven unscheduled interruptions due to software problems (mostly ASP) and seven because of hardware. There were 17 days of uninterrupted service and 28 days in which more than 22 hours of TSO service was provided.

	May 1976	June 1976
Hours available (JPS and TSO)	401.2	441.56
Daily average of hours available	12.9	14.7
Hours of unattended TSO available	234.9	244.1
Daily average of hours of unattended TSO	7.5	8.1
Hours lost due to failures	71.9	4.6
Hardware	67.0	2.0
Software	4.9	2.6
Number of failures	34	18
Mean time between failures	18.7 hours	38.1 hours
Mean time to repair	127 minutes	15.3 minutes
Actual/scheduled uptime	90%	99%

TSO Statistics for June 1976

These statistics are for M.I.T. usage of TSO only.

Number of LOGON sessions	8,957
Total CONNECT time	5,221 hours
Average CONNECT time per session	35 minutes
Total CPU time	19 hours
Average CPU time per session	0.13 minutes
Total disk I/O operations	9.6 million
Average disk I/O operations per session	1,073
Total teleprocessing I/O operations	2.5 million
Average teleprocessing I/O operations per session	284
Total K byte minutes of interactive memory use	0.74 million
Average K byte minutes per session	83

JOB PROCESSING SYSTEM STATISTICS FOR JUNE 1976

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	21,542	
Number of M.I.T. job steps processed on the 168	73,922	
Total CPU time used	98.2	hours
Average CPU time used per job	0.31	minutes
*Total run time used (estimated)	259	hours
Average run time per job	0.72	minutes
Total number of I/O operations by type		
CTC	11.0	million
Tape	2.9	million
Disk	15.1	million
Average number of I/O operations per job		
CTC	581	
Tape	1,120	
Disk	781	
Total memory minutes used	3.6	million K bytes
Average memory minutes used per job	184	K bytes
Total number of lines printed	30.0	million
Average number of lines printed per job	1,411	
Total number of cards read	6.2	million
Average number of cards read per job	290	
Total number of cards punched	0.62	million
Average number of cards punched per job	448	
Total number of Remote Job I/O units	9.7	thousand
Average number of RJ I/O units per job	569	
Total number of setup jobs	4,636	
Total number of setup tapes	9,819	
Total number of setup disks	7,157	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP or MFT system.

TURN-AROUND STATISTICS FOR JUNE 1976

The following table reports turn-around figures for the Job Processing System for the month of June, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		9,352	11,193	315	21,767
Average Turn-around		43 mins.	15 mins.	25 mins.	121 mins.
Percentage of Jobs Returned	25%	8 mins.	4 mins.	5 mins.	6 mins.
	50%	23 mins.	9 mins.	10 mins.	12 mins.
	75%	55 mins.	14 mins.	24 mins.	29 mins.
	95%	149 mins.	44 mins.	44 mins.	239 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 17,042 non-setup jobs was . . . . . .101 mins.
The average turn-around for 4,725 setup jobs was . . . . . .195 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 100K	2,695	137 mins.	11 mins.
102 - 200K	9,939	87 mins.	11 mins.
202 - 300K	5,488	130 mins.	12 mins.
302 - 400K	2,388	165 mins.	14 mins.
402 - 500K	651	200 mins.	24 mins.

MULTICS

I/O Daemon Accounting to Change

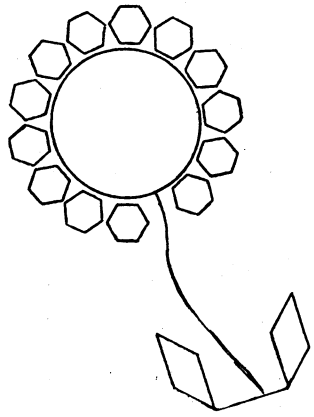

In order to make charges more representtative of the actual costs of printing, the units used for charging purposes by the I/O Daemon will be changed from blocks (700 bits) to lines on July 1. The price per unit will remain the same. On average, this should represent no price change to the user, but segments with longer lines will be less expensive while segments with short lines will be more expensive than at present.

A New Look for Multics Info Segments

M.I.T. and Honeywell are now making extensive revisions to the Multics system's info segments. New info files are being added for all commands which now lack them; we are updating and reformating others and deleting those which are obsolete.

This is being done in several stages. As a first step, we have replaced the contents of the directory >doc>is, which contains the "standard" info segments, with segments in the new format.

This format, which is much more terse than most of the old info segments, makes it easier for the user to get quick and concise information on a command and its arguments, but does not contain complete information and explanation. To supplement these info segments, we will re-establish the on-line MPM; until this is done, the old info segments will be left on line in >doc>old_info_segments (ois).

We encourage you to use the new info segments and would appreciate feedback about their usefulness (submit comments to the Multics consultant via mail to consultant.consult (c.c)). For further information on these new info segments, type "help pending_changes".

Multics Operating Statistics for May 1976

	April	May	% Change
System hours up	612	534	-12.7
Two CPU hours	525	388	-26.1
MTBF (hours)	16:05	15:42	-2.4
CPU charged (hours)	263	269	+2.3
Memory units (million)	12.3	12.0	-2.4
Print lines (million)	10.6	9.8	-7.6
Terminal hours	15,753	15,134	-3.9
Logins	15,387	13,921	-9.5
Absentee jobs	3,272	2,659	-18.7
Disk usage (page-months)	128,263	123,826	-3.5
Projects	225	220	-2.2
Users	1,963	1,959	-0.2

Back Up Your 2314 Disk Packs with Fast Dump Restore

The Fast Dump Restore procedure (FDR) can be used to back up your private 2314 disk packs to tape and to restore them from tape to disk. FDR is more efficient than the IBM utility IEHDASDR.

For instance, to dump the 2314 disk pack 234000 to standard-labeled tape 001234, you would use the following job stream:

```
      Job parameter statements
      . . .
/*SETUP UNIT=2314,ID=234000,A=XYZ
/*SETUP UNIT=TAPE9,ID=(001234,RING,SAVE,SL),A=ABC
//       EXEC FDR,DISK=234000,TAPE=001234
/*
```

To restore the dumped disk pack from tape 001234 to a new 2314 disk pack, 234111, you would use this job stream:

```
      Job parameter statements
      . . .
/*SETUP UNIT=2314,ID=234111,A=DEF
/*SETUP UNIT=TAPE9,ID=(001234,NORING,SAVE,SL),A=ABC
//     EXEC FDR,PRM=N,DISK=234111,TAPE=001234
/*
```

FDR also has special functions which allow you to dump and restore specific datasets. A new Center memo, AP-16, soon to be published, details these functions.

Cataloged Procedures for Allocating and Cataloging Datasets

The Center now provides two cataloged procedures, ALLOC and REALLOC, for use in allocating disk space and cataloging on-line datasets more easily.

To allocate and catalog a new empty sequential dataset with the name U.M12345.P67890.USER.FORT, with a blocksize of 3120 and a maximum of 450 card images, use the following job stream:

```
      Job parameter statements (see Center memo OS-9)
      . . .
//  EXEC ALLOC,DS='U.M12345.P67890.USER.FORT',TRACKS=3,TYPE=DECK
/*
```

To allocate and catalog a new empty sequential dataset with the above name and characteristics, forcibly deleting any cataloged dataset existing with that name, use the following job stream:

```
      Job parameter statements (see Center memo OS-9)
      . . .
//  EXEC REALLOC,DS='U.M12345.P67890.USER.FORT',TRACKS=3,TYPE=DECK
/*
```

ALLOC and REALLOC can also be used to allocate partitioned datasets. For further information, see the new Center memo OS-2, which will be published soon.

Check Out the Checkout Compiler

For debugging PL/I programs on OS, the IBM Checkout compiler is a very useful tool. It accepts the same language as the IBM Optimizing compiler, but it gives more detailed error messages and does extensive monitoring for potential error conditions (such as checking for uninitialized variables). Although it can be used in the batch environment, it is most useful on TSO, where patches can be made in mid-execution and then selectively saved in a permanent copy of the source program. The recently revised Center memo OS-13 describes the Checkout compiler and gives many examples of its time-sharing facilities. Center memo PP-3, a new revision of which will soon be available, describes its use in the batch environment.

On TSO an error condition causes control to be passed to the terminal, giving the programmer many options. The programmer can also gain control at will by pressing the "attention" (or "quit" or "break") button which causes execution to be interrupted and control to be passed to the programmer. Typing "help" whenever control is received will generate a message indicating the current phase of processing and a list of the actions valid at the given moment. Typing "help <item>" will give syntax and function descriptions of any specified action. Valid actions include Checkout subcommands or PL/I statements which are to be executed immediately or inserted into the code as a patch (to be executed when encountered in the normal logic flow of the program).

A typical sequence would be as follows. After executing nine statements ("9X"), the ERROR condition is raised, because the variable J is being used (in statement 115) before it has been assigned a value. Control passes to the terminal with the prompting message "?", indicating the execution phase of processing. At this point the programmer wishes to print the source statement ("list 115") and the attributes of the variable J ("reference (j) x,a"). Then the programmer makes a temporary assignment to J and requests that execution be resumed at the point where it was originally interrupted ("go to 115").

```
IEN1004I 9X 'J' HAS NO VALUE.
    AT 115 IN G_C_D
IEN1187A 'ERROR' RAISED.
? list 115
       115 J=J+1;
? ref (j) x,a
**** J AUTOMATIC ALIGNED BINARY FIXED(15,0)
  115, 220
? j=0; go to 115
```

The sequence above uses "immediate PL/I"; that is, the value of the variable is assigned immediately but the code of the program is not altered. It is also possible to make a permanent patch to the program source by inserting a statement. For example,

```
? 105 j=0; go to 0
```

will insert the statement "J=0;" as statement 105 (or replace an existing statement with the same number) and then begin execution anew from the beginning of the program ("go to 0").

The above is a very superficial introduction to the use of the Checkout compiler in the time-sharing environment; for a more comprehensive treatment, see Center memo OS-13.

Speakeasy Removed

Because of apparent lack of interest in this product and other operational considerations, we removed the Speakeasy system from OS TSO on June 7, 1976. Anyone interested in supporting possible reinstallation should contact Arthur Anger in Room 39-469, Telephone 253-7044, or Linda Desmond in Room 39-475, Telephone 253-6322.

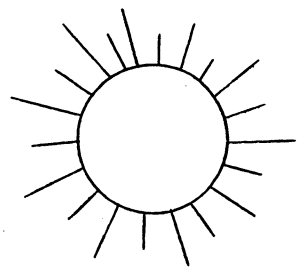

S/370 Performance in May 1976

The first week in May was a disaster as far as hardware reliability was concerned. We had twelve crashes caused by hardware problems that week; one of them resulted in fifteen hours of down time. After that things did get better, but not much; we had only ten days with no crashes and our time lost due to hardware failures (67 hours) exceeded the total for the previous four months. And all this during the end-of-term rush!

	April 1976	May 1976
Hours available (JPS and TSO)	401.4	401.2
Daily average of hours available	13.4	12.9
Hours of unattended TSO available	268.4	234.9
Daily average of hours of unattended TSO	8.9	7.5
Hours lost due to failures	12.9	71.9
Hardware	8.1	67.0
Software	4.8	4.9
Number of failures	24	34
Mean time between failures	27.9 hours	18.7 hours
Mean time to repair	32 minutes	127 minutes
Actual/scheduled uptime	98%	90%

TSO Statistics for May 1976

These statistics are for M.I.T. usage of TSO only.

Number of LOGON sessions	6,796	
Total CONNECT time	3,988	hours
Average CONNECT time per session	35	minutes
Total CPU time	16	hours
Average CPU time per session	0.14	minutes
Total disk I/O operations	7.4	million
Average disk I/O operations per session	1,087	
Total teleprocessing I/O operations	1.8	million
Average teleprocessing I/O operations per session	265	
Total K byte minutes of interactive memory use	0.56	million
Average K byte minutes per session	82	

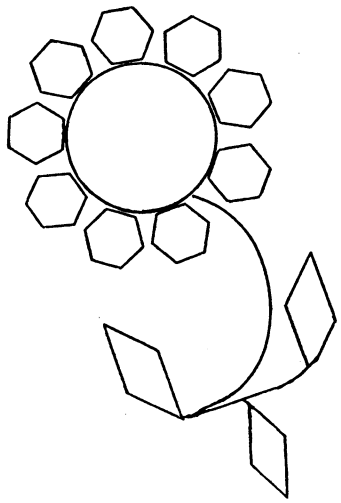

JOB PROCESSING SYSTEM STATISTICS FOR MAY 1976

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	19,465	
Number of M.I.T. job steps processed on the 168	67,321	
Total CPU time used	95.7	hours
Average CPU time used per job	0.33	minutes
*Total run time used (estimated)	234	hours
Average run time per job	0.72	minutes
Total number of I/O operations by type		
CTC	8.9	million
Tape	2.0	million
Disk	13.5	million
Average number of I/O operations per job		
CTC	514	
Tape	908	
Disk	767	
Total memory minutes used	3.4	million K bytes
Average memory minutes used per job	190	K bytes
Total number of lines printed	26.1	million
Average number of lines printed per job	1,360	
Total number of cards read	5.6	million
Average number of cards read per job	288	
Total number of cards punched	0.48	million
Average number of cards punched per job	420	
Total number of Remote Job I/O units	1.3	thousand
Average number of RJ I/O units per job	193	
Total number of setup jobs	3,651	
Total number of setup tapes	8,034	
Total number of setup disks	5,657	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP or MFT system.

TURN-AROUND STATISTICS FOR MAY 1976

The following table reports turn-around figures for the Job Processing System for the month of May, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		7,433	10,916	647	19,731
Average Turn-around		60 mins.	18 mins.	14 mins.	112 mins.
Percentage of Jobs Returned	25%	10 mins.	5 mins.	4 mins.	6 mins.
	50%	28 mins.	10 mins.	8 mins.	13 mins.
	75%	74 mins.	26 mins.	13 mins.	43 mins.
	95%	223 mins.	59 mins.	29 mins.	254 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 15,985 non-setup jobs was . . . . . 94 mins.
The average turn-around for 3,746 setup jobs was . . . . . .188 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 100K	2,052	124 mins.	12 mins.
102 - 200K	9,099	92 mins.	12 mins.
202 - 300K	5,265	111 mins.	13 mins.
302 - 400K	2,000	129 mins.	23 mins.
402 - 500K	775	174 mins.	24 mins.

(Reading Room continued)

M.I.T. - Rotch Library

Census Data: A Selected List of Guides and Information Sources. 1976.

Northwestern University

Computation of Bicubic-Spline Surfaces for Irregularly-Spaced Data. 1975.

Pennsylvania State University

Accurate Fourier Analysis of Arbitrary Functions Defined by Discrete Data Values. 1975.

Prentice-Hall, Incorporated

APL: The Language and Its Usage. 1975.

SofTech

AED Installation and Maintenance Manual. 1971.

University of California, Berkeley

A Program to Compute the Condition Numbers of Matrix Eigenvalues Without Computing Eigenvectors. 1975.

An Improved Shift Strategy for the QR-Algorithm for Real Hessenberg Matrices. 1976.

Observational Research in User-Computer Interaction. 1975.

John Wiley & Sons

Microprocessors and Microcomputers. 1976.

New Message Facility Installed

The interactive message commands (send_message, etc.) have been updated to use ring 1 Person.mbx mailboxes. Their calling sequences have not changed, except that 1) there is a new send_message_express command (to send a message only if the user is accepting messages); 2) a new control argument (-last, to print the last message received) has been added to print_messages; and 3) the imf_state command is no longer available.

For compatibility with the old message facility, print_messages prints the contents of Person.con_msgs as well as the contents of the user's mailbox. We recommend, however, that you delete your Person.con_msgs segment, since it is no longer necessary.

Entries in the old facility will be available temporarily with the names old_message_facility$send_message, etc.

Type "help new_message_facility" (nmf) for a more complete discussion of the new message facility.

* * *

Services for RDMS Users

RDMS users now have their own on-line consultant and trouble report facilities. Both the rdms_online_consultant (rolc) and rdms_trouble_report (rtr) commands are used like the standard Multics commands (olc and tr), which are documented in the IML (Center memo MS-13) available in the Publications Office, Room 39-484, Telephone 253-6325.

These commands are only available to registered RDMS users.

AML Utility for FORTRAN Programmers

Here's good news for the user who must import multi-subprogram FORTRAN programs onto the Multics system. A new Author-Maintained Library utility called split_fortran_subprograms.teco (alias sfs.teco) should make the process less tedious.

Teco (text editor and corrector), a powerful character-oriented text editor installed in the AML, can be used to create complex macros which perform "automatic" editing of segments. The sfs.teco utility is a macro which was created using Teco; it separates imported FORTRAN subprograms residing in one segment into single subprogram segments as required by the Multics FORTRAN compiler. The macro does this by searching from the bottom upwards for valid "END" statements and separating the subprograms at those points; it creates a sub-directory called "teco_fortran" in which to store segments containing the separated subprograms. The segments thus created are named after the subprograms (i.e., they are called (subprogram_name).FORTRAN). For instance, a subprogram called "FUNCTION SUB1" is stored in a segment called "sub1.FORTRAN" after being separated. Note that the contents of the source and output segments of the programs consist of upper-case text.

To use the sfs.teco macro to separate a multi-subprogram segment, you invoke the Teco text editor; read the external macro "sfs.teco" into the buffer; and execute it, supplying the name of the multi-subprogram segment as requested. After you return to command level, you should check the contents of the sub-directory teco_fortran. Note that the main program and the block data program(s) without names will be found in segments called mainpx.FORTRAN, where x is a numeric character. You may then use the AML command c360f to change all the separated subprograms from upper-case to lower-case text, so that they will be acceptable to the Multics FORTRAN compiler. The subprogram can then be compiled.

For more information about the sfs.teco macro, including an annotated example, see the write-up in the April additions to the AML. Details about the Teco editor and the c360f command are also available in the Multics AML documentation, which may be obtained in the Publications Office (Room 39-484).

Multics Operating Statistics for April 1976

	March	April	% Change
System hours up	636	612	-3.8
Two CPU hours	580	525	-9.6
MTBF (hours)	18:41	16:05	-13.9
CPU charged (hours)	342	263	-23.1
Memory units (million)	15.4	12.3	-20.2
Print lines (million)	14.4	10.6	-26.4
Terminal hours	17,670	15,753	-10.8
Logins	16,514	15,387	-6.8
Absentee jobs	4,678	3,272	-30.1
Disk usage (page-months)	163,210	128,263	-21.4
Projects	233	225	-3.4
Users	1,996	1,963	-1.7

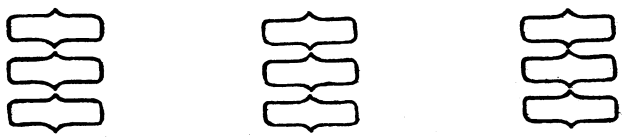

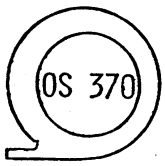

Edition 5 of IMSL Installed

Edition 5 of the IMSL library was installed on the IBM 370 on May 10. While most users of the library should not notice any difference, there have been substantial changes in the library, with the addition of fifty-six new subroutines and the deletion of thirty-five old subroutines from the previous edition. In addition, sixty subroutines from Edition 4 have been improved.

Highlights of the changes include:

1. A subroutine for the real generalized eigenvalue/eigenvector problem $Ax = \lambda Bx$.
2. A major revision to the cubic-spline subroutines in Chapters D and I.
3. An improved space economizer and a new high-accuracy linear-equation solver for indefinite symmetric matrices.
4. Six new matrix storage-mode conversion subroutines.
5. A major revision of the polynomial-root finder based on Laguerre's method.
6. A major revision of the quadratic equation solvers.
7. A new nonlinear least-squares subroutine that uses a finite Levenberg-Marquardt algorithm.
8. A new subroutine to find the minimum of a function of N variables using a quasi-Newton method.
9. ARIMA time-series model analysis including model identification and maximum-likelihood estimation.
10. A set of subroutines allowing a complete factor analysis, including estimation and rotation options.
11. The Furnival and Wilson "leaps and bounds" algorithm for model selection through determination of a number of best regression subsets.
12. Kelvin functions.

The new cubic-spline subroutines in Chapters D and I have standard interfaces for the spline coefficients to facilitate changing from one subroutine to another (e.g., from interpolation to least squares). In addition, the SOR iterative method, which was subject to convergence failure, has been replaced by a direct method.

Users of the subroutine BECORO should note that the length of the parameter TEMP has been increased.

The random-number generators GGU1 and GGU2 have been replaced by GGUB and GGU3. The new subroutines, which have been subjected to theoretical and empirical tests, are more efficient. All IMSL subroutines which use random numbers now call GGUB rather than GGU1.

ZXMARQ has been revised by IMSL advisor Professor Kenneth Brown. The new version, ZXSSQ, has an improved iterative scheme for choosing the Levenberg-Marquardt parameter and an option to force descent.

ZXPOWL has been replaced by a quasi-Newton method that is named ZXMIN. This subroutine has better convergence properties than ZXPOWL for most problems and is able to handle a wider range of problems.

Other subroutines which have been replaced are EQZHSF/EQZHF, EQZISF/EQZIF, EQZVAF/EQZVEF, IQHERU, LEQT1S, ZPOLYR, and ZQUADR.

Anyone encountering difficulty with the new version of the library should consult with Mark Froimowitz, Room 39-465, Telephone 253-7729. Obsolete subroutines will also be available, if necessary, from him.

S/370 Performance in April 1976

Although we had several more interruptions in service this month than during March, the total down time was still less than 2 percent of the scheduled up time. The longest interruption was a three-hour delay caused by a software problem. We did have fifteen days with no interruptions and a large number of days when the system came up early.

	March 1976	April 1976
Hours available (JPS and TSO)	404.8	401.4
Daily average of hours available	13.0	13.4
Hours of unattended TSO available	277.8	268.4
Daily average of hours of unattended TSO	9.0	8.9
Hours lost due to failures	1.0	12.9
Hardware	0.5	8.1
Software	0.5	4.8
Number of failures	6	24
Mean time between failures	113.8 hours	27.9 hours
Mean time to repair	10 minutes	32 minutes
Actual/scheduled uptime	99.8%	98%

* * *

TSO Statistics for April 1976

These statistics are for M.I.T. usage of TSO only.

Number of LOGON sessions	8,866
Total CONNECT time	5,229 hours
Average CONNECT time per session	35 minutes
Total CPU time	21 hours
Average CPU time per session	0.15 minutes
Total disk I/O operations	11.3 million
Average disk I/O operations per session	1,279
Total teleprocessing I/O operations	2.5 million
Average teleprocessing I/O operations per session	283
Total K byte minutes of interactive memory use	0.84 million
Average K byte minutes per session	94

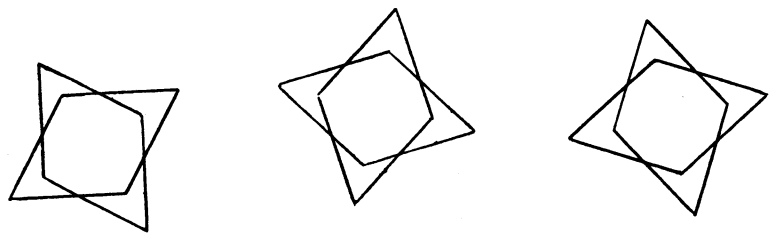

Job Processing System Statistics for April 1976

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	21,507	
Number of M.I.T. job steps processed on the 168	70,797	
Total CPU time used	88.2	hours
Average CPU time used per job	0.27	minutes
*Total run time used (estimated)	232	hours
Average run time per job	0.65	minutes
Total number of I/O operations by type		
CTC	9.0	million
Tape	2.1	million
Disk	14.0	million
Average number of I/O operations per job		
CTC	466	
Tape	1,082	
Disk	710	
Total memory minutes used	3.3	million K bytes
Average memory minutes used per job	169	K bytes
Total number of lines printed	27.7	million
Average number of lines printed per job	1,306	
Total number of cards read	6.6	million
Average number of cards read per job	310	
Total number of cards punched	0.48	million
Average number of cards punched per job	379	
Total number of setup jobs	3,518	
Total number of setup tapes	7,625	
Total number of setup disks	5,493	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP or MFT system.

TURN-AROUND STATISTICS FOR APRIL 1976

The following table reports turn-around figures for the Job Processing System for the month of April, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		9,531	11,032	520	21,815
Average Turn-around		44 mins.	17 mins.	8 mins.	108 mins.
Percentage of Jobs Returned	25%	7 mins.	4 mins.	4 mins.	6 mins.
	50%	22 mins.	9 mins.	9 mins.	12 mins.
	75%	44 mins.	14 mins.	13 mins.	29 mins.
	95%	164 mins.	44 mins.	30 mins.	224 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 18,194 non-setup jobs was . . . . . 98 mins.
The average turn-around for 3,621 setup jobs was . . . . . .200 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 100K	2,408	163 mins.	11 mins.
102 - 200K	9,726	72 mins.	11 mins.
202 - 300K	5,742	126 mins.	12 mins.
302 - 400K	2,557	91 mins.	12 mins.
402 - 500K	805	146 mins.	14 mins.

MULTICS OPERATING STATISTICS FOR MARCH

	February	March	% Change
System hours up	540	636	+17.6
Two CPU hours	419	580	+38.5
MTBF (hours)	11:29	18:41	+62.6
CPU charged (hours)	245	342	+39.9
Memory units (million)	10.9	15.4	+41.7
Print lines (million)	10.3	14.4	+39.4
Terminal hours	13,169	17,670	+34.2
Logins	14,167	16,514	+16.6
Absentee jobs	3,309	4,678	+41.4
Disk usage (page-months)	128,215	163,210	+27.3
Projects	235	233	-0.9
Users	2,100	1,996	-5.0

AN INVITATION: TRY SPEAKEASY IN MAY

A new subsystem is available on TSO. Called Speakeasy, this subsystem greatly simplifies access to computing power for many types of calculation. When used as a desk calculator, it performs not only the usual arithmetic operations, but can also:

handle arrays of numbers automatically
perform logical tests
evaluate standard mathematical functions
estimate derivatives from a table
compute frequency distributions and correlations
compute inverses and eigenvalues

These operations may also be collected in stored programs for convenient reuse. Here are some sample Speakeasy statements demonstrating operations on an array (the colon-underscore is an input-prompting signal):

```
:_x=array(2,4:integers(9,16))
:_sumrows(x)
SUMROWS(X) (A 2 COMPONENT ARRAY)
  42  58
:_sumcols(x)
SUMCOLS(X) (A 4 COMPONENT ARRAY)
  22  24  26  28
:_average(x)
AVERAGE(X) = 12.5
:_variance(x)
VARIANCE(X) = 6
:_transpose(x)
TRANSPOSE(X) (A 4 by 2 ARRAY)
  9   13
  10  14
  11  15
  12  16
:_
```

(continued on page 10)

(Speakeasy continued)

Except for a command which undoes the effects of the immediately preceding computation, Speakeasy has no features which cannot be found in other widely used languages. Its advantage is that it includes a wide range of capabilities and makes many of them fairly convenient to use. Furthermore, while Speakeasy provides enough prompting messages and explanatory material to help the beginner find the right way to do simple things, it has enough options and flexibility to allow construction of quite sophisticated computational structures. FORTRAN programmers may find it useful as a convenient interactive control for their existing routines which may be executable with little adaptation.

We suggest you try Speakeasy at your terminal; you need only type:

```
logon userid size(300)
speakez
```

You can sample its arithmetic capabilities by typing a few simple expressions and assignments, test its reference information by typing HELP, and be guided through the subsystem's fundamentals by typing TUTORIAL. For more expansive (and less expensive) documentation, you may then consult the manuals and articles on reserve in the IPC Reading Room, 39-430. We would appreciate receiving your comments about the uses you find for Speakeasy during May. Please address comments or questions to Linda Desmond, Room 39-475, Telephone 253-6322.

Tektronix Libraries to be Revised

The Center will install Release 3.1 of the Tektronix Terminal Control Systems at the end of May. To provide compatibility with this new terminal control system, changes will be made to some of the Advanced Graphing II and CalComp Preview routines.

The exact installation date for new versions of these Tektronix libraries (SYS5.TLIB4002.LOAD, SYS5.TLIB4010.LOAD, SYS5.TLIBAG2.LOAD, SYS5.TLIBPLOT.LOAD) will be advertised in tailsheet and TSO broadcast messages.

Release 3.1 and Advanced Graphing II routines should be used as documented in the vendor documents, copies of which are available in the IPC Reading Room, 39-430. The only user-written software which might be incompatible with the new libraries is that which was written at a low level; for example, programs relying on the Release 2.1 common block TKTRNX will have to be revised because this has been changed. Contact the Programming Assistance and Information (PAI) Office (Room 39-217, Telephone 253-4114) if you have questions.

The change to the TCAM translate tables mentioned in the January 1976 Bulletin, which will require a revision of the Tektronix Terminal Control System libraries, has been postponed indefinitely.

New User Identification for SC-4020 Plots

In the past, the project and programmer numbers to be used in charging for SC-4020 plots were obtained from the PARM field of the spooling procedure SC4020. Unfortunately, when program decks were passed from one user to another, the project and programmer numbers often were not changed, causing plots to be charged to the wrong account.

To solve this problem, a system change has been made so that the user's identification will be obtained from the /*MITID or the /*HUID statement, instead of from the spool step. This change is transparent to all users and does not require any changes on your part. If you have any difficulty with SC-4020 plotting charges, please check with the consultants in the PAI Office (Room 39-219).

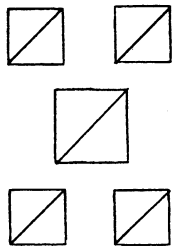

More Errors Corrected in SPSSH Version 6

On April 5, 1976, three errors in SPSSH were corrected; they were in the following areas:

1) The significance of chi-square: It was calculated for the wrong tail of the distribution if the degrees of freedom exceeded 60 and the degrees of freedom were greater than the value of chi-square. This error may have had an impact on extracting functions or loadings, as in DISCRIMINANT, CANCORR, and FACTOR procedures.

2) DISCRIMINANT with error 1812: An incorrect variable specification resulted in a print loop for the error.

3) ARCHIVE: Files containing value labels were not always read properly, but generated error 1273.

Questions may be directed to Suzanne Chen (Telephone 253-7769).

S/370 Performance in March

Performance reliability in March was exceptional. An unscheduled IPL was required only four different days and with six total failures we only lost one hour of uptime. Result: 99.8 per cent uptime!

	February 1976	March 1976
Hours available (JPS and TSO)	332.0	404.8
Daily average of hours available	11.4	13.0
Hours of unattended TSO available	247.0	277.8
Daily average of hours of unattended TSO	8.5	9.0
Hours lost due to failures	12.1	1.0
Hardware	10.0	0.5
Software	2.1	0.5
Number of failures	26	6
Mean time between failures	22.3 hours	113.8 hours
Mean time to repair	28 minutes	10 minutes
Actual/scheduled uptime	98%	99.8%

TSO Statistics for March

Number of LOGON sessions	8,267
Total CONNECT time	5,323 hours
Average CONNECT time per session	39 minutes
Total CPU time	22 hours
Average CPU time per session	0.16 minutes
Total disk I/O operations	11.1 million
Average disk I/O operations per session	1,337
Total teleprocessing I/O operations	2.3 million
Average teleprocessing I/O operations per session	288
Total K byte minutes of interactive memory use	0.81 million
Average K byte minutes per session	97

Job Processing System Statistics for March

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	20,969	
Number of M.I.T. job steps processed on the 168	72,510	
Total CPU time used	85.8	hours
Average CPU time used per job	0.27	minutes
*Total run time used (estimated)	239	hours
Average run time per job	0.68	minutes
Total number of I/O operations by type		
CTC	10.9	million
Tape	1.8	million
Disk	15.0	million
Average number of I/O operations per job		
CTC	582	
Tape	771	
Disk	789	
Total memory minutes used	3.5	million K bytes
Average memory minutes used per job	182	K bytes
Total number of lines printed	29.0	million
Average number of lines printed per job	1,401	
Total number of cards read	6.5	million
Average number of cards read per job	309	
Total number of cards punched	0.49	million
Average number of cards punched per job	365	
Total number of setup jobs	4,220	
Total number of setup tapes	3,391	
Total number of setup disks	2,046	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP or MFT system.

TURN-AROUND STATISTICS FOR MARCH

The following table reports turn-around figures for the Job Processing System for the month of March, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		9,395	10,610	450	21,393
Average Turn-around		37 mins.	15 mins.	10 mins.	116 mins.
Percentage of Jobs Returned	25%	7 mins.	4 mins.	4 mins.	6 mins.
	50%	15 mins.	9 mins.	9 mins.	12 mins.
	75%	43 mins.	14 mins.	13 mins.	29 mins.
	95%	134 mins.	44 mins.	44 mins.	270 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 17,114 non-setup jobs was . . . . . 95 mins.
The average turn-around for 4,279 setup jobs was . . . . . .196 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 100K	2,727	92 mins.	11 mins.
102 - 200K	9,374	94 mins.	11 mins.
202 - 300K	5,827	129 mins.	12 mins.
302 - 400K	2,291	133 mins.	13 mins.
402 - 500K	644	163 mins.	22 mins.

News for Users of BASIC Random Files

For efficiency, a revised I/O system which handles random string and number files in a new format has been installed for the BASIC compiler. This I/O system will automatically convert any old-format file it encounters to the new format (if access and quota permit), print a message, and continue processing.

If a file cannot be converted automatically, the message "unable to convert old format file filename to new format" will be printed and the BASIC program will be aborted. At that point the user may check access, etc., and type

```
convert_old_basic_file filename
```

which will convert the file to the new format. This command can also be used prior to processing for files known to be of the old format.

IMSL on Multics

Edition 4 of the International Mathematical and Statistical Library (IMSL) is now available on Multics. For complete information on how to use this library of subroutines, see Center memo AP-9.

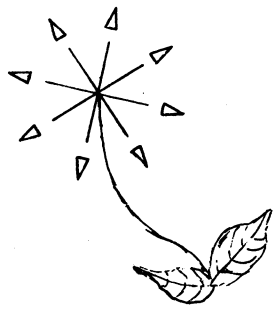

Multics Operating Statistics for February

	January	February	% Change
System hours up	572	540	-5.5
Two CPU hours	482	419	-13.1
MTBF (hours)	16:19	11:29	-29.6
CPU charged (hours)	331	245	-26.0
Memory units (million)	12.8	10.9	-15.3
Print lines (million)	10.8	10.3	-4.1
Terminal hours	14,102	13,169	-6.6
Logins	13,695	14,167	+3.4
Absentee jobs	3,609	3,309	-8.3
Disk usage (page-months)	142,243	128,215	-9.9
Projects	233	235	+0.9
Users	1,977	2,100	+6.2

Improvements in 370/168 Data Lines for TSO

We have upgraded the TSO data line distribution to accommodate the many new 300-baud terminals acquired by users over the past few months. There are now thirty-two lines numbered consecutively from (617) 258-7511 through 258-7542.

Improvements included installing additional lines and assigning consecutive extension numbers. Previously, users experienced problems either because all lines were busy or because they were unsure of what extensions to try since the numbers were not consecutive.

We hope these changes will correct both problems.

Plotting Rate Changes Coming

CalComp

Because a great deal of time is required to set up CalComp India ink (drafting pen) plots, we have decided to require a minimum charge of $10.00 per job on these plots. The charge for changes in color or pen size, however, has been reduced from $5.00 to $2.00 per change. (The $30.00 per hour of plot time rate remains the same.)

We'd like to remind you that liquid (India) ink plots, though great for professional reproduction, are difficult to set up and require constant operator attention; we encourage you to develop and debug graphs using ballpoint pen. The same plot tape can then be used to replot only the final graph in India ink, avoiding the expense and difficulty of repeated India ink plotting.

Quick debugging of plots on TSO is also possible using CalComp previewing on Tektronix terminals like the one available for public use in the Reading Room (39-430). We encourage you to read Center memo AP-29 for more information on previewing CalComp plots.

Please note that since drafting pens run out of ink after about 45 minutes, we require you to make special arrangements if you think a plot will take longer than this, so that an appropriate "break point" to change pens can be found.

SC-4020

Our vendor has increased its per-frame rate for SC-4020 work and we must correspondingly increase ours. The new rates are:

Output	Price per frame
paper only	$0.75
film only	$0.75
both	$1.00

See the March 10, 1976 Addendum to Center memo GI-7 for details on both CalComp and SC-4020 rate changes. Please note that these changes will become effective on April 1, not on March 1 as stated in the addendum.

Version 1.4 of TSO Utilities Installed

New versions of the TSO Data Utilities -- COPY, FORMAT, MERGE and LIST -- were recently installed. The only M.I.T. modification that should affect users is that NONUM is now the default for the COPY command. Please see the consultants in the Programming Assistance and Information (PAI) Office (Room 39-219) if you have any problems with these utilities.

Trial Run for SPEAKEASY

The Center will soon install the Speakeasy subsystem on TSO for several months of evaluation. Speakeasy offers the colloquial style of BASIC, the array-handling operations and interactive control of APL, and the program-library accessibility of FORTRAN or PL/I. Developed originally for use in physics, it is also convenient for solving problems in mathematics, engineering, and statistics. Descriptive literature is on reserve in the Reading Room (39-430). Watch system announcements or type "help news" for installation date and execution instructions.

Corrections to Errors in SPSSH Version 6

Known errors in Version 6.01 of SPSS (listed in the January Bulletin) were corrected on February 24. The errors occurred under some circumstances in processing control statements ADD CASES, ANOVA, AGGREGATE, ARCHIVE, BREAKDOWN, CONDESCRIPTIVE, CROSSTABS, DISCRIMINANT, INPUT FORMAT-BINARY, IF or SELECT IF, REGRESSION, and ONEWAY. Detailed descriptions of the errors can be found in the SPSS Newsletter No. 9 which is posted in the Programming Assistance and Information (PAI) Office (Room 39-219) and at the East Campus facility.

Correction of these errors necessitates no change in the usage of SPSS cataloged procedures. Users should continue to specify VERSION=V600 on the EXEC statement when using the second edition of the SPSS manual.

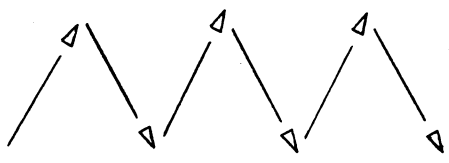

Recommendation for XTAPEDMP

To reduce charges for the general purpose tape-dumping program XTAPEDMP, the Center recommends use of the BYPASS command instead of SPACE (or SKIP) for bypassing files. BYPASS requires only one input/output operation per file bypassed, while SPACE (or SKIP) requires one input/output operation per block. This can result in large input/output costs when many files are bypassed. Note that BYPASS does not record the block count for bypassed files.

For example, to get a summary dump (the first eighty characters of each physical block) of the fourteenth file of a nine-track standard-labeled tape, use the following format:

```
// 'MYNAME'
/*MITID USER=(M12345,P67890,,,mypass)
/*MAIN LINES=1
/*SETUP UNIT=TAPE9,ID=(123456,NORING,SAVE,BLP),A=XXX
//  EXEC XTAPEDMP,VOL=123456,UNIT=TAPE9
//SYSIN  DD  *
  BYPASS 39 FILES
  PRINT   3 FILES SUMMARY
/*
```

Note that the fourteenth standard-labeled file is equivalent to BLP files 40, 41, and 42.

S/370 Performance in February

We had a large number of interruptions in February, and, although most resulted in very short periods of down time, on two days we were down for extended periods because of hardware problems. Incidentally, of the twenty-six interruptions, six were caused by environmental problems (either power or chilled water failures).

	January 1976		February 1976	
Hours available (JPS and TSO)	374.1		332.0	
Daily average of hours available	12.5		11.4	
Hours of unattended TSO available	262.0		247.0	
Daily average of hours of unattended TSO	8.5		8.5	
Hours lost due to failures	38.3		12.1	
Hardware	37.3		10.0	
Software	1.0		2.1	
Number of failures	16		26	
Mean time between failures	39.8	hours	22.3	hours
Mean time to repair	144	minutes	28	minutes
Actual/scheduled uptime	94%		98%	

TSO Statistics for February

Number of LOGON sessions	6,737	
Total CONNECT time	4,111	hours
Average CONNECT time per session	37	minutes
Total CPU time	15	hours
Average CPU time per session	0.14	minutes
Total disk I/O operations	8.1	million
Average disk I/O operations per session	1,210	
Total teleprocessing I/O operations	1.9	million
Average teleprocessing I/O operations per session	286	
Total K byte minutes of interactive memory use	0.58	million
Average K byte minutes per session	87	

JOB PROCESSING SYSTEM STATISTICS FOR FEBRUARY

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	16,659	
Number of M.I.T. job steps processed on the 168	57,292	
Total CPU time used	71.0	hours
Average CPU time used per job	0.30	minutes
*Total run time used (estimated)	207	hours
Average run time per job	0.74	minutes
Total number of I/O operations by type		
CTC	7.1	million
Tape	2.6	million
Disk	12.7	million
Average number of I/O operations per job		
CTC	502	
Tape	1,282	
Disk	877	
Total memory minutes used	2.9	million K bytes
Average memory minutes used per job	199	K bytes
Total number of lines printed	22.3	million
Average number of lines printed per job	1,392	
Total number of cards read	4.9	million
Average number of cards read per job	292	
Total number of cards punched	0.36	million
Average number of cards punched per job	347	
Total number of setup jobs	3,587	
Total number of setup tapes	2,851	
Total number of setup disks	1,950	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP or MFT system.

TURN-AROUND STATISTICS FOR FEBRUARY

The following table reports turn-around figures for the Job Processing System for the month of February, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		6,971	8,832	387	16,964
Average Turn-around		48 mins.	16 mins.	20 mins.	104 mins.
Percentage of Jobs Returned	25%	8 mins.	5 mins.	5 mins.	6 mins.
	50%	24 mins.	10 mins.	9 mins.	13 mins.
	75%	57 mins.	26 mins.	15 mins.	42 mins.
	95%	209 mins.	59 mins.	44 mins.	239 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 13,353 non-setup jobs was . . . . . 85 mins.
The average turn-around for 3,611 setup jobs was . . . . . .173 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 100K	2,535	90 mins.	13 mins.
102 - 200K	7,146	80 mins.	12 mins.
202 - 300K	4,495	124 mins.	13 mins.
302 - 400K	1,831	123 mins.	14 mins.
402 - 500K	475	154 mins.	22 mins.

(Reading Room continued)

PDP-11 Peripherals Handbook. 1975.

LSI-11, PDP-11/03 Processor Handbook. 1975.

PDP-11/45 Processor Handbook. 1973.

PDP-11, RSTS/E System Manager's Guide. 1973.

Honeywell Information Systems, Incorporated

Multics Programmers' Manual - Reference Guide. 1975.

IBM Research Center

Multiprocessing for Real-Time Applications. 1975.

Structured Programming with and without GO TO Statements. 1975.

General Register Assignment in Presence of Data Flow. 1975.

Introduction to Data Translation. 1975.

On-Line Generation of Terminological Digests in Language Translation. 1975.

Computer Aided Drafting at IBM Research. 1975.

Technology for Storing Extendible Arrays. 1975.

On the Self-Contained Modelling of DB/DC Systems. 1975.

A Note on the Power of Recursion and Equality Versus That of Parallelism. 1975.

Granularity of Locks and Degrees of Consistency in a Shared Data Base. 1975.

Statistical Analysis of Transaction Processing in a Data Base System. 1975.

Institute of Electrical and Electronics Engineers

Proceedings of the First National Conference on Software Engineering. 1975.

An Algorithm for Exponentiation to Fractional Powers and Rooting. 1974.

M.I.T. - Laboratory of Architecture and Planning

Consistent System Programmer's Introduction. 1974.

M.I.T. - Project MAC

MACLISP Reference Manual, Parts 1 and 2 of six parts. 1975.

Triangle Universities Computation Center

List of Current TUCC/NCECS Documentation. 1975.

Report for Fiscal Year, 1974-1975.

University of Illinois

A General Method for Evaluation of Functions and Computations in a Digital Computer. 1975.

Numerical Errors in Sparse Linear Equations. 1975.

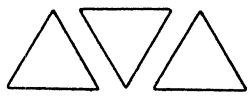

MULTICS OPERATING STATISTICS FOR JANUARY

	December	January	% Change
System hours up	674	572	-15.2
Two CPU hours	642	482	-24.9
MTBF (hours)	24:04	16:19	-32.1
CPU charged (hours)	249	331	+32.7
Memory units (million)	13.1	12.8	-2.5
Print lines (million)	11.2	10.8	-3.8
Terminal hours	15,760	14,102	-10.5
Logins	13,711	13,695	-0.1
Absentee jobs	3,364	3,609	+7.3
Disk usage (page-months)	155,921	137,350	-11.9
Projects	222	233	+5.0
Users	1,792	1,977	+10.3

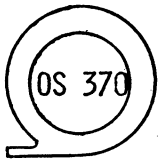

NEW CLIST FOR FORTRAN DEFINE FILE DATASET ALLOCATION

A new CLIST, ALLOCDA, allows TSO FORTRAN programmers to allocate space for new DEFINE FILE datasets and to call a program to write the required skeleton records. If you've been accustomed to the previously circuitous method of using direct access files on TSO, we think you'll be delighted with the short cut provided by this CLIST.

The ALLOCDA CLIST is used as follows:

```
allocda dsname recno(n) recsize(m) [list]
```

"dsname" should be the dataset name in the short form (i.e., without the prefix "U.M12345.P67890.", since this logon id prefix will be assumed).

The value given for "n" should be the maximum number of records in the file.

The value given for "m" should be the maximum size of a record measured in bytes.

The optional operand "list" prints the CLIST as it is executed.

For a program that contains the statement

```
DEFINE FILE 8(100,80,L,INDX)
```

you could create and format the dataset defi.data for direct access processing using the new CLIST as follows:

```
allocda defi.data recno(100) recsize(80)
```

Then logical unit 8 is allocated to the dataset:

```
alloc fi(ft08f001) da(defi.data)
```

and the compiled program is executed:

```
loadgo myprog fortlib
```

Refer to Center memo PP-15 for more information on DEFINE FILE datasets.

S/370 Performance in January

System availability was excellent for the first twenty days of January. A disk problem on the twenty-third caused us considerable grief, however, and most of our fifteen hardware problems were related to that. Happily, only one software problem occurred during the month. Because of a number of engineering changes, an additional twenty hours of maintenance were scheduled.

	December 1975	January 1976
Hours available (JPS and TSO)	333.2	374.1
Daily average of hours available	11.5	12.5
Hours of unattended TSO available	361.6	262.0
Daily average of hours of unattended TSO	11.7	8.5
Hours lost due to failures	5.1	38.3
Hardware	0.5	37.3
Software	4.6	1.0
Number of failures	16	16
Mean time between failures	36.5 hours	39.8 hours
Mean time to repair	16 minutes	144 minutes
Actual/scheduled uptime	99%	94%

TSO Statistics for January

Number of LOGON sessions	6,498	
Total CONNECT time	3,942	hours
Average CONNECT time per session	36	minutes
Total CPU time	15	hours
Average CPU time per session	0.14	minutes
Total disk I/O operations	7.6	million
Average disk I/O operations per session	1,170	
Total teleprocessing I/O operations	1.7	million
Average teleprocessing I/O operations per session	263	
Total K byte minutes of interactive memory use	.0.55	million
Average K byte minutes per session	85	

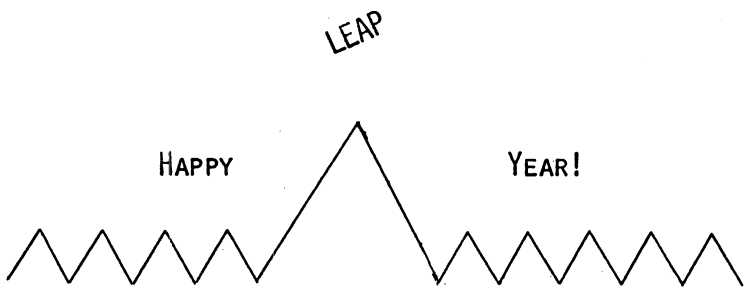

JOB PROCESSING SYSTEM STATISTICS FOR JANUARY

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	17,327	
Number of M.I.T. job steps processed on the 168	33,355	
Total CPU time used	57.7	hours
Average CPU time used per job	0.22	minutes
*Total run time used (estimated)	193	hours
Average run time per job	0.67	minutes
Total number of I/O operations by type		
CTC	8.6	million
Tape	1.9	million
Disk	13.0	million
Average number of I/O operations per job		
CTC	568	
Tape	1,090	
Disk	831	
Total memory minutes used	2.7	million K bytes
Average memory minutes used per job	176	K bytes
Total number of lines printed	21.7	million
Average number of lines printed per job	1,274	
Total number of cards read	5.1	million
Average number of cards read per job	297	
Total number of cards punched	0.33	million
Average number of cards punched per job	313	
Total number of Remote Job I/O units	4.8	thousand
Average number of RJ I/O units per job	444	
Total number of setup jobs	3,303	
Total number of setup tapes	7,102	
Total number of setup disks	5,045	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP or MFT system.

TURN-AROUND STATISTICS FOR JANUARY

The following table reports turn-around figures for the Job Processing System for the month of January, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		7,217	9,530	390	17,664
Average Turn-around		43 mins.	14 mins.	14 mins.	78 mins.
Percentage of Jobs Returned	25%	8 mins.	4 mins.	5 mins.	6 mins.
	50%	23 mins.	9 mins.	10 mins.	11 mins.
	75%	55 mins.	14 mins.	14 mins.	29 mins.
	95%	150 mins.	45 mins.	44 mins.	149 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 14,316 non-setup jobs was . . . . . 64 mins.
The average turn-around for 3,348 setup jobs was . . . . . .138 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 100K	2,194	110 mins.	11 mins.
102 - 200K	7,449	62 mins.	11 mins.
202 - 300K	5,428	72 mins.	11 mins.
302 - 400K	1,616	93 mins.	13 mins.
402 - 500K	353	119 mins.	21 mins.

IMSL Notes

In order to correct various bugs, changes were made in mid-January to the following International Mathematical and Statistical Library (IMSL) routines on the IBM system: BDCOU1, DVOGER, FTMAXL, GGU3, ICSFKU, ICSFKV, LUDAPB, MDTPOS/MDPOS, USWBSM, and VPOLYF. In addition, users of NKS1 should check their results for reasonableness since some difficulties have been reported. If you require further information about these routines, call Mark Froimowitz (Telephone 253-7729).

December's Top Twenty Programs on the 370: Two Viewpoints

In December, 33,545 different job steps and 379 different programs were executed by M.I.T. users of the 370/168, using 61 hours of CPU time. The lists below show the programs which were most heavily used in terms of frequency of execution and CPU usage. Users' programs are identified by an asterisk (*).

By Frequency of Execution

			% of 33,545 Job Steps
1.	IGIFORT	FORTRAN G1 Compiler	13.3%
2.	IEWLDRGO	Loader	10.3%
3.	PGM=*.DD	Users' Programs	9.4%
4.	IEWL	Linkage Editor	8.3%
5.	IEBPTPCH	IBM Utility Program	7.6%
6.	WATFIV	Waterloo's FORTRAN Compiler	7.6%
7.	IEFBR14	Null Program	7.3%
8.	IEBGENER	IBM Utility Program	5.2%
9.	IEL0AA	PL/I Optimizing Compiler	2.5%
10.	IGHRC000	Sort/Merge	1.9%
11.	CALCMP00	CalComp	1.5%
12.	IEKAA00	FORTRAN H Compiler	1.5%
13.	SPSSV600	Statistical Package for the Social Ssciences-Version 6	1.4%
14.	DYN2	DYNAMO	0.9%
15.	TSP	Time Series Processor	0.8%
16.	IEBCOPY	IBM Utility Program	0.8%
17.	MAPDISK	Disk Contents Lister	0.7%
18.	SYSJCL	*	0.6%
19.	IEBUPDTE	IBM Utility Program	0.6%
20.	IEHLIST	IBM Utility Program	0.5%

By CPU Usage

			% of Total CPU Usage for All Programs
1.	PGM=*.DD	Users' Programs	36.0%
2.	IEWLDRGO	Loader	22.5%
3.	WATFIV	Waterloo's FORTRAN Compiler	8.7%
4.	IGIFORT	FORTRAN G1 Compiler	5.1%
5.	IEL0AA	PL/I Optimizing Compiler	2.2%
6.	TEMPNAME	Users' Programs	1.5%
7.	OVERLORD	*	1.3%
8.	MAXITSP	Time Series Processor	1.3%
9.	DJLEXEC	MPSX (Mathematical Programming System Extended)	1.0%
10.	IEKAA00	FORTRAN H Compiler	0.9%
11.	SPSSV600	Statistical Package for the Social Sciences-Version 6	0.9%
12.	DAILY	IPC Billing Program	0.8%
13.	IEWL	Linkage Editor	0.8%
14.	SPLPRT	Multics Spool Print	0.8%
15.	MAPDISK	Disk Contents Lister	0.7%
16.	LBRREC	*	0.7%
17.	DATATEXT	Statistical Analysis Program	0.6%
18.	WS	*	0.6%
19.	IEBPTPCH	IBM Utility Program	0.5%
20.	DYN2	DYNAMO	0.5%

S/370 Performance in December

During the month of December, the Center was "closed" for about three days for the holidays, reducing the number of Job Processing System hours. During that time we ran unattended TSO service (seventy consecutive hours in one stretch) and this naturally increased the average number of hours for this service. Reliability still left a lot to be desired, particularly in the software area, but we did have nineteen days with no interruption in service.

	November 1975	December 1975
Hours available (JPS and TSO)	355.9	333.2
Daily average of hours available	12.2	11.5
Hours of unattended TSO available	279.0	361.6
Daily average of hours of unattended TSO	9.3	11.7
Hours lost due to failures	12.8	5.1
Hardware	6.5	.5
Software	6.3	4.6
Number of failures	19	19
Mean time between failures	33.4 hours	36.5 hours
Mean time to repair	40 minutes	16 minutes

TSO Statistics for December

Number of LOGON sessions	6,412
Total CONNECT time	3,958 hours
Average CONNECT time per session	37 minutes
Total CPU time	16 hours
Average CPU time per session	0.15 minutes
Total disk I/O operations	8.7 million
Average disk I/O operations per session	1,351
Total teleprocessing I/O operations	1.8 million
Average teleprocessing I/O operations per session	291
Total K byte minutes of interactive memory use	.0.63 million
Average K byte minutes per session	99

FROM ALL OF US AT THE CENTER

TO

ALL OF YOU . . .

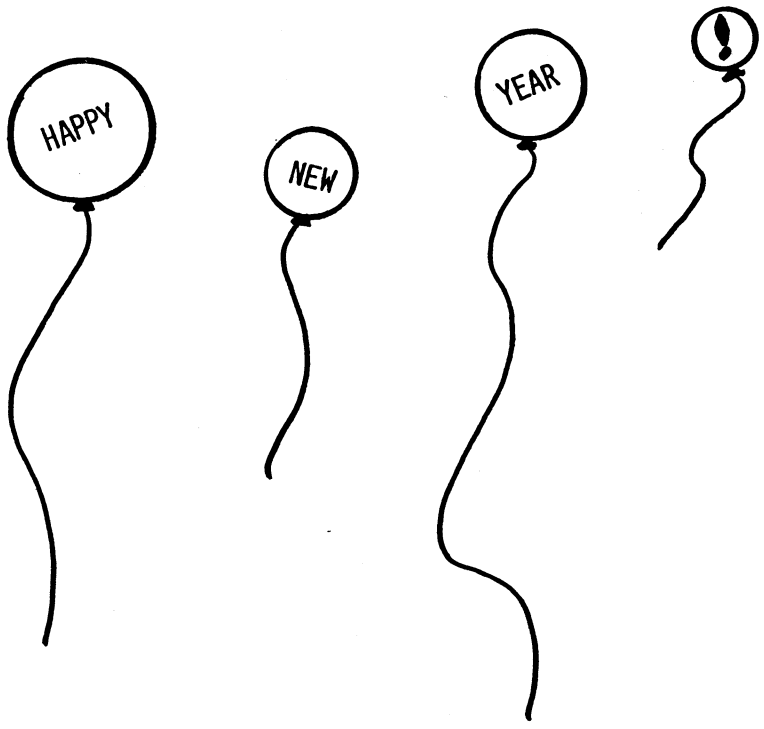

JOB PROCESSING SYSTEM STATISTICS FOR DECEMBER

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	17,058	
Number of M.I.T. job steps processed on the 168	33,545	
Total CPU time used	61.3	hours
Average CPU time used per job	0.24	minutes
*Total run time used (estimated)	201	hours
Average run time per job	0.71	minutes
Total number of I/O operations by type		
CTC	8.1	million
Tape	2.6	million
Disk	13.5	million
Average number of I/O operations per job		
CTC	537	
Tape	1,419	
Disk	875	
Total memory minutes used	2.9	million K bytes
Average memory minutes used per job	188	K bytes
Total number of lines printed	23.3	million
Average number of lines printed per job	1,384	
Total number of cards read	4.9	million
Average number of cards read per job	289	
Total number of cards punched	0.39	million
Average number of cards punched per job	342	
Total number of Remote Job I/O units	3.1	thousand
Average number of RJ I/O units per job	437	
Total number of setup jobs	3,198	
Total number of setup tapes	2,779	
Total number of setup disks	1,525	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP or MFT system.

TURN-AROUND STATISTICS FOR DECEMBER

The following table reports turn-around figures for the Job Processing System for the month of December, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		7,045	8,866	520	17,056
Average Turn-around		44 mins.	16 mins.	12 mins.	117 mins.
Percentage of Jobs Returned	25%	7 mins.	5 mins.	4 mins.	6 mins.
	50%	22 mins.	10 mins.	9 mins.	12 mins.
	75%	44 mins.	15 mins.	13 mins.	29 mins.
	95%	59 mins.	59 mins.	45 mins.	209 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 13,858 non-setup jobs was . . . . . 87 mins.
The average turn-around for 3,198 setup jobs was . . . . . .251 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 100K	2,185	137 mins.	11 mins.
102 - 200K	8,425	80 mins.	11 mins.
202 - 300K	4,038	151 mins.	12 mins.
302 - 400K	1,358	134 mins.	12 mins.
402 - 500K	476	219 mins.	13 mins.

(Bibliographies continued)

If you'd like to know where to find information on specified topics related to System/360 in IBM's Systems Reference Library publications, you may want to check the IBM System/360 Operating System Master Index (GC28-6644, $3.40). This manual lists topics alphabetically with references to System Reference Library publications where they are discussed. The Reading Room has copies of this index for both reference and circulation.

The Bibliography of Data Processing Techniques (GF20-8172, $0.25) lists a selection of IBM publications which contain descriptions of data processing techniques-- for instance, generally useful methods and applications, studies in analysis and design, management of a data processing department, etc. Each entry provides a title, order number, current status, and an abstract. Publications are listed both by topic with order numbers indicated, and in order-number sequence with titles given. This manual is available for borrowing in the Reading Room.

Once you've obtained some manuals and started reading, you may find another IBM reference manual useful for obtaining definitions of many terms used in the data processing field. This is the Data Processing Glossary (GC20-1699, $1.70), which contains over 2,000 definitions from the American National Standard Vocabulary for Information Processing, published by the American National Standards Institute, Inc. (ANSI) as well as a proposed new American National Standard Vocabulary, and several glossaries published previously by IBM.

These bibliographies should give you an idea about the range of IBM publications. Please note that the Publications Office generally stocks only those manuals relevant to the hardware and software currently installed at the Center. However, individual copies of any IBM manual, except publications relating to licensed program products which are not on our system, can be ordered through this office.

S/370 Performance in November

The month started out well-- nine days with no unscheduled interruptions of service; after that problems started. Only eight of the remaining twenty-one days were trouble-free and several days had more than one interruption. We feel we are carrying some of the problems into December, though we're working hard to correct them.

Incidentally, since unattended TSO service started in March, we've had only six interruptions in the service, and it appears to be working well.

	October 1975	November 1975
Hours available (JPS and TSO)	410.8	355.9
Daily average of hours available	13.2	12.2
Hours of unattended TSO available	263.0	279.0
Daily average of hours of unattended TSO	8.5	9.3
Hours lost due to failures	16.6	12.8
Hardware	13.0	6.5
Software	3.6	6.3
Number of failures	16	19
Mean time between failures	25.7 hours	33.4 hours
Mean time to repair	62 minutes	40 minutes

TSO Statistics for November

Number of LOGON sessions	7,142	
Total CONNECT time	4,312	hours
Average CONNECT time per session	36	minutes
Total CPU time	19	hours
Average CPU time per session	0.16	minutes
Total disk I/O operations	8.7	million
Average disk I/O operations per session	1,221	
Total teleprocessing I/O operations	2.1	million
Average teleprocessing I/O operations per session	297	
Total K byte minutes of interactive memory use	0.71	million
Average K byte minutes per session	99	

JOB PROCESSING SYSTEM STATISTICS FOR NOVEMBER

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	17,871	
Number of M.I.T. job steps processed on the 168	34,304	
Total CPU time used	56.7	hours
Average CPU time used per job	0.21	minutes
*Total run time used (estimated)	194	hours
Average run time per job	0.65	minutes
Total number of I/O operations by type		
CTC	8.7	million
Tape	1.9	million
Disk	13.7	million
Average number of I/O operations per job		
CTC	547	
Tape	946	
Disk	848	
Total memory minutes used	2.7	million K bytes
Average memory minutes used per job	165	K bytes
Total number of lines printed	23.2	million
Average number of lines printed per job	1,317	
Total number of cards read	5.3	million
Average number of cards read per job	296	
Total number of cards punched	0.4	million
Average number of cards punched per job	324	
Total number of Remote Job I/O units	3.1	thousand
Average number of RJ I/O units per job	310	
Total number of setup jobs	3,488	
Total number of setup tapes	2,746	
Total number of setup disks	1,694	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP or MFT system.

TURN-AROUND STATISTICS FOR NOVEMBER

The following table reports turn-around figures for the Job Processing System for the month of November, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		8,035	8,823	376	17,872
Average Turn-around		37 mins.	14 mins.	16 mins.	94 mins.
Percentage of Jobs Returned	25%	7 mins.	4 mins.	4 mins.	6 mins.
	50%	21 mins.	9 mins.	9 mins.	11 mins.
	75%	42 mins.	14 mins.	14 mins.	28 mins.
	95%	134 mins.	44 mins.	45 mins.	179 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 14,384 non-setup jobs was . . . . . 74 mins.
The average turn-around for 3,488 setup jobs was . . . . . 175 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 100K	2,020	124 mins.	11 mins.
102 - 200K	9,247	72 mins.	11 mins.
202 - 300K	4,368	104 mins.	12 mins.
302 - 400K	1,367	99 mins.	13 mins.
402 - 500K	472	204 mins.	23 mins.

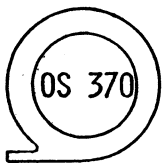

How to Protect Your Datasets from Yourself!

M.I.T.'s on-line dataset access control mechanism (also known as the dataset protection facility), implemented in February of 1974, was designed to prevent unauthorized access to user datasets, and also to preclude malicious or accidental deletion of system or user datasets. An extra feature of this facility is that you can use it to prevent yourself from accidentally deleting your own important datasets.

Setting OWNER(R) (owner read only) access to your dataset will allow you to read it. But if you should accidentally attempt to DELETE (or RENAME) such a dataset, the access control mechanism would prevent this.

For example, to prevent datasets with the names U.M12345.P67890.RESEARCH.DATA and U.M12345.P67890.ANALYSIS.LOAD from being accidentally deleted or renamed, you could enter the following TSO command:

```
permit (research.data,analysis.load) OWNER(R)
```

Alternatively, you could submit a batch job of the following form:

```
//  NAME
/*MITID  USER=(M12345.P67890,,,password)
//  EXEC  PERMIT
//SYSIN  DD  *
 PERMIT U.M12345.P67890.RESEARCH.DATA OWNER(R)
 PERMIT U.M12345.P67890.ANALYSIS.LOAD OWNER(R)
```

Note that you can always reset the access to read and write (OWNER(RW)) to prepare for deletion of any datasets. You must also reset the access to read and write before attempting to update (or write into) a dataset. See Center memo OS-39 for more information on how to set access on datasets.

After setting access on a dataset, you should then issue the SPACE command (see Center memo OS-22) or check the USERMAP to be sure the access control attributes have actually been properly set.

FORTRAN normally opens a dataset for INOUT or OUTIN (input and output) depending on whether a READ or WRITE statement, respectively, is encountered first. In a FORTRAN program which will read a dataset with OWNER(R) access, you must override this normal method of opening a dataset. Opening such a dataset for INOUT violates the dataset security system, causing your job to abend with system error 913, even if you do not try to write into the dataset. This is because INOUT allows the dataset to be opened for input and then later used for output without reopening it. Thus write (W) access is required, even if you don't want to write into the dataset. Since your OWNER(R) dataset does not allow you write access, you should open the dataset for INPUT only by including the LABEL parameter in JCL or using the ATTRIB and ALLOCate commands in order to read these datasets, as follows:

```
//G.FT16F001  DD  DSN=U.M12345.P67890.RESEARCH.DATA,DISP=SHR,
//          LABEL=(,,,IN)
```

or

```
attrib readonly input
alloc fi(ft16f001) da(research.data) using(readonly)
```

FORTRAN hintbug 9, copies of which are available in the PAI Office (Room 39-219), gives more information about this use of the LABEL parameter. See the TSO Command Language Reference manual (GC28-6732-4) for details about the ATTRIB and ALLOC commands.

Three New Volumes Replace OS 360/370 Bibliography

IBM's OS 360/370 Bibliography (GA22-6822) is now obsolete and has been replaced by three new volumes which you can purchase in the Publications Office (Room 39-483, Telephone (617) 253-6325). GC20-0001, 370 Bibliography, is available for $13.50, while the 360 Bibliography (GC20-0360) costs $15.00. GC20-0370, the Applications Bibliography, is being sold for $5.20.

SPSS Version 6 Corrected

Version 6 of SPSS has been corrected for known errors in BREAKDOWN, FACTOR, CROSSTABS, and ANOVA in addition to the errors listed in No. 8 of the SPSS Newsletter.

Users of the Second Edition of the SPSS manual should specify Version 6 on the EXEC statement as follows in order to employ the new facilities (see the June, 1975 issue of the Bulletin):

```
// EXEC SPSSX,VERSION=V600
```

Version 4.00 of SPSSH will be deleted at the end of this year, but version 5.00 will remain.

More Texts for Beginners

If you are interested in learning more about the IBM operating system at the Center, here are three IBM manuals which may help you.

Introduction to IBM Data Processing Systems (GC20-1684, $3.00) is a student text designed to present concepts and principals essential to a basic knowledge of computers. The manual covers in basic terms such topics as data representation, storage devices, CPU functions, I/O devices, operating systems, etc. General principals are illustrated with references to specific IBM systems.

Once you have learned some basic computer principles from that manual, you may want to read GC28-6534, IBM System/360 Operating System Introduction ($1.90) which describes operating systems in general before introducing the 360 system. This publication describes the purpose, design objectives, organization, function, and application of this operating system; how it was influenced by previous systems; and its compatibility with System/370.

Another student text, Introduction to IBM System/360 Architecture (GC20-1667, $1.10), describes such features of System/360 architecture as channels, automatic interrupts, and general purpose registers. Storage addressing, instruction formats, data formats, and the various types of arithmetic operations are also discussed. Questions and exercises are provided for review.

S/370 Performance in October

There were sixteen service interruptions during the month of October, eleven hardware related and five software related. Of the eleven hardware-caused problems, six, which occurred over a five-day period, were traced to a poorly connected cable. One 3330 problem caused a three-and-one-half hour delay in service early in the month.

	September 1975	October 1975
Hours available (JPS and TSO)	392.6	410.75
Daily average of hours available	13.1	13.2
Hours of unattended TSO available	231.1	263.0
Daily average of hours of unattended TSO	7.7	8.5
Hours lost due to failures	1.6	16.6
Hardware	0.5	13.0
Software	1.1	3.6
Number of failures	11	16
Mean time between failures	56.7 hours	25.7 hours
Mean time to repair	5 minutes	62 minutes

M.I.T. USAGE OF IBM 370/168
OCTOBER 1974 - OCTOBER 1975

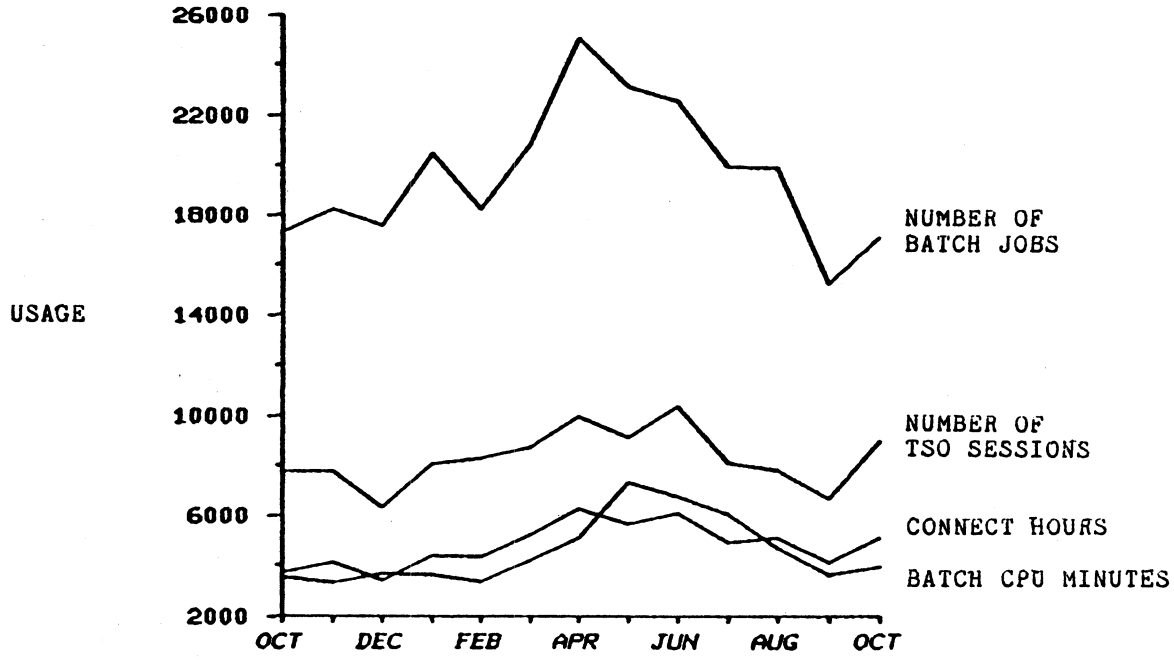

TSO Statistics for October

Number of LOGON sessions	8,944	
Total CONNECT time	5,060	hours
Average CONNECT time per session	34	minutes
Total CPU time	18	hours
Average CPU time per session	0.12	minutes
Total disk I/O operations	9.1	million
Average disk I/O operations per session	1,019	
Total teleprocessing I/O operations	2.4	million
Average teleprocessing I/O operations per session	269	
Total K byte minutes of interactive memory use	0.69	million
Average K byte minutes per session	77	

JOB PROCESSING SYSTEM STATISTICS FOR OCTOBER

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	17,074	
Number of M.I.T. job steps processed on the 168	33,793	
Total CPU time used	65.5	hours
Average CPU time used per job	0.26	minutes
*Total run time used (estimated)	209	hours
Average run time per job	0.73	minutes
Total number of I/O operations by type		
CTC	8.3	million
Tape	2.1	million
Disk	14.4	million
Average number of I/O operations per job		
CTC	553	
Tape	1,111	
Disk	939	
Total memory minutes used	3.0	million K bytes
Average memory minutes used per job	199	K bytes
Total number of lines printed	22.3	million
Average number of lines printed per job	1,328	
Total number of cards read	5.2	million
Average number of cards read per job	305	
Total number of cards punched	0.44	million
Average number of cards punched per job	403	
Total number of Remote Job I/O units	4.3	thousand
Average number of RJ I/O units per job	2,135	
Total number of setup jobs	3,455	
Total number of setup tapes	2,700	
Total number of setup disks	1,795	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP or MFT system.

TURN-AROUND STATISTICS FOR OCTOBER

The following table reports turn-around figures for the Job Processing System for the month of October, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		7,607	8,578	307	17,068
Average Turn-around		32 mins.	17 mins.	12 mins.	97 mins.
Percentage of Jobs Returned	25%	6 mins.	4 mins.	5 mins.	5 mins.
	50%	13 mins.	9 mins.	10 mins.	11 mins.
	75%	29 mins.	14 mins.	25 mins.	27 mins.
	95%	104 mins.	44 mins.	59 mins.	135 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 13,616 non-setup jobs was . . . . . 82 mins.
The average turn-around for 3,452 setup jobs was . . . . . 158 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 100K	2,171	128 mins.	11 mins.
102 - 200K	7,925	78 mins.	10 mins.
202 - 300K	4,356	83 mins.	11 mins.
302 - 400K	1,666	100 mins.	12 mins.
402 - 500K	505	333 mins.	14 mins.

"HELP" HAS BECOME MORE HELPFUL

A new version of the "help" command, installed in September, offers you a more convenient and controlled way to obtain information about Multics commands, subroutines, and other features.

The new help command, like the old one, pauses at intervals to allow the user to request further information or to exit from the help file. Paragraphs are no longer delimited with the printer-on (006) character; instead, "help" looks for a double empty line signalling the end of a paragraph. Info segments written in the old format are compatible with the new command, however, since special code checks for 00o characters and uses any that are found as pause delimiters. The command now asks "more help" after every paragraph.

"help" recognizes a phrase followed by a colon found at the beginning of a paragraph as a section title, and uses that title in the "more help" query. An example of the new query format is:

```
7 lines titled "Arguments" follow.  More help?
```

When you receive the query "more help", your choice of answers is no longer limited to "yes" or "no." You may still, of course, answer "yes" to request printing of the next section of the info file, or "no" to indicate that you are finished with the segment, but you now have the following other choices:

quit	indicates that you are finished with "help".
rest	requests printing of the rest of the info segment.
search xxx yyy	requests "help" to search for the indicated strings.
section SSS	requests "help" to search for the section with the given title.
skip	requests "help" to skip the next block.
title	requests "help" to print titles only of the sections comprising the rest of the segment.

When searching for both sections and strings is requested, section search is done first.

If you find that a help file is so long that you do not wish to wait while it is printed at your terminal, you may dprint the info file after determining its location by typing "help xx -he", where xx is the title of the help file. For instance, if you type "help help -he", you receive the answer

```
>doc>info>help.info 07/23/75
```

indicating the segment in which the help file is located and the date of its most recent modification.

Info Segment Format

Each info segment begins with a brief title line, beginning with the date of the file's last modification. For command or subroutine descriptions, this line gives the name(s) of the program, including abbreviations.

Command descriptions include sections on "Function," "Syntax," and (if applicable) "Arguments" and "Control Arguments." Other sections titled "Notes," "Examples," "Changes," or "References" will be included if necessary.

The "Function" section gives a one- or two-line description of what the command does. A sample invocation of the command is provided in the "Syntax" section. The "Arguments" and "Control Arguments" sections give one-line descriptions of each argument or control argument.

"Function," "Syntax," and "Arguments" sections are provided in subroutine info segments. ("Syntax" and "Arguments" sections are provided for each entry point.)

Most help files are still in the old format; however, they are being rewritten as they come up for revision. Many features of the new help command are applicable to old-format write-ups, but some, such as those which require section titles, cannot be used. We hope that Multics info files will be completely converted to the new format in the not too distant future, enabling users to take advantage of all features of the new help command.

MULTICS OPERATING STATISTICS FOR SEPTEMBER

	August	September	% Change
System hours up	610	677	+10.9
Two CPU hours	562	573	+2.0
MTBF (hours)	27:44	27:04	-2.4
CPU charged (hours)	173	178	+3.1
Memory units (million)	8.0	7.5	-6.9
Print lines (million)	7.7	8.4	+8.6
Terminal hours	13,432	13,209	-1.7
Logins	10,437	12,746	+22.1
Absentee jobs	2,333	2,031	-12.9
Disk usage (page-months)	93,943	111,594	+18.8
Projects	207	212	+2.4
Users	1,311	1,626	+24.0

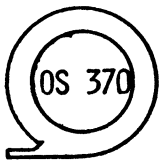

P-STAT DOCUMENTS RELOCATED

Because it has been infrequently used, the P-STAT documentation (Manual, Primer, and Examples) has been renamed and moved to a publicly accessible disk pack.

You can use the following job format to have one of these documents printed (where { } indicates that you should choose one of the options given):

```
//  'name',REGION=65K,CLASS=A
/*MITID  userparm
/*SRI    level
/*MAIN   LINES= {10} for {Manual  }
               {1 }     {Primer  }
               {4 }     {Examples}
/*SETUP  UNIT=2314,ID=234130,A=WYH
//  EXEC PGM=IEBGENER
//SYSIN  DD  DUMMY
//SYSPRINT  DD  SYSOUT=A
//SYSUT1  DD  DSNAME=PV.PSTAT.{MANUAL},DISP=SHR,UNIT=2314,
                            {PRIMER }
                            {EXAMPL }
//    VOL=SER=234130
//SYSUT2  DD  SYSOUT=A
/*
```

The cataloged procedures PSMANUAL, PSPRIMER, and PSEXAMPL are no longer available.

MORE HELP FOR NEW COMPUTER USERS

This month we'll continue our series of articles on IBM student texts available in the Publications Office with descriptions of three sets of manuals designed for self-study or in-house training courses. If you're just introducing yourself to the Center's IBM computer, or know someone who is, these may be useful.

The first three manuals comprise an introductory course on "Computing Systems Fundamentals," designed for beginners in data processing. These three manuals -- Overview Text (GR29-0280, $3.75), Techniques Text (GR29-0281, $6.95), and Notebook (GR29-0282, $11.45) should give the reader a basic knowledge of the way a computer works and how it can be used to solve problems. Decision tables, flowcharts, and appropriate documentation for problem solutions are discussed.

(continued on page 11)

(Help, continued)

"Fundamentals of Programming" (SR29-0019, $2.40) is recommended for those planning to take a programming language course, who have a knowledge of the basic information contained in the first set of manuals. This text teaches the reader how to handle typical data processing problems, covering flowcharting, decision making, analysis, testing, and other programming requirements. Also available are an Advisor's Guide (SR29-0015, $0.68) and Advisor's Guide Supplement (SR29-0018, $0.10) for those who are advising students using the text.

A similar set of text and advisor's manuals present an "Introduction to System/360." The Text (GR29-0256, $10.00) and Notebook (GR29-0257, $1.30) are designed to give students an understanding of the architecture, components, and programming systems of IBM's System/360. This course assumes a background in the material covered by "Computing Systems Fundamentals," described above. The Advisor's Guide mentioned in the "Fundamentals of Programming" write-up is also applicable to this course; an Advisor's Guide Supplement for "Introduction to System/360" (GR29-0255) is available for $0.10.

Any of these manuals may be purchased in the Publications Office, Room 39-483, Telephone 253-6325.

S/370 Performance in September

There were twenty-one days in September with no unscheduled interruptions in operation. We had a total of eleven crashes, three hardware-related, four software, all of them minor; service was available for 87 percent of the 720 hours of the month.

	August 1975		September 1975	
Hours available (Batch and TSO)	408.4		392.6	
Daily average of hours available	13.2		13.1	
Hours of unattended TSO available	273.7		231.1	
Daily average of hours of unattended TSO	8.8		7.7	
Hours lost due to failures	8.6		1.6	
Hardware	5.2		0.5	
Software	3.4		1.1	
Number of failures	19		11	
Mean time between failures	35.8	hours	56.7	hours
Mean time to repair	27	minutes	5	minutes

TSO Statistics for September

Number of LOGON sessions	6,670	
Total CONNECT time	4,092	hours
Average CONNECT time per session	37	minutes
Total CPU time	16	hours
Average CPU time per session	0.14	minutes
Total disk I/O operations	7.9	million
Average disk I/O operations per session	1,190	
Total teleprocessing I/O operations	1.8	million
Average teleprocessing I/O operations per session	270	
Total K byte minutes of interactive memory use	0.55	million
Average K byte minutes per session	83	

JOB PROCESSING SYSTEM STATISTICS FOR SEPTEMBER

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	15,216	
Number of M.I.T. job steps processed on the 168	32,674	
Total CPU time used	60.3	hours
Average CPU time used per job	0.26	minutes
*Total run time used (estimated)	207	hours
Average run time per job	0.82	minutes
Total number of I/O operations by type		
CTC	7.1	million
Tape	2.3	million
Disk	14.7	million
Average number of I/O operations per job		
CTC	527	
Tape	1,333	
Disk	1,004	
Total memory minutes used	3.0	million K bytes
Average memory minutes used per job	219	K bytes
Total number of lines printed	20.2	million
Average number of lines printed per job	1,348	
Total number of cards read	5.1	million
Average number of cards read per job	333	
Total number of cards punched	0.57	million
Average number of cards punched per job	483	
Total number of Remote Job I/O units	2.9	thousand
Average number of RJ I/O units per job	483	
Total number of setup jobs	3,385	
Total number of setup tapes	2,549	
Total number of setup disks	1,929	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP or MFT system.

TURN-AROUND STATISTICS FOR SEPTEMBER

The following table reports turn-around figures for the Job Processing System for the month of September, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		6,555	7,523	442	15,216
Average Turn-around		33 mins.	15 mins.	18 mins.	116 mins.
Percentage of Jobs Returned	25%	6 mins.	4 mins.	5 mins.	5 mins.
	50%	12 mins.	9 mins.	10 mins.	11 mins.
	75%	28 mins.	14 mins.	14 mins.	26 mins.
	95%	33 mins.	15 mins.	18 mins.	345 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 11,831 non-setup jobs was . . . . . 97 mins.
The average turn-around for 3,385 setup jobs was . . . . . .182 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 100K	2,036	163 mins.	11 mins.
102 - 200K	7,418	83 mins.	10 mins.
202 - 300K	3,221	133 mins.	11 mins.
302 - 400K	1,635	120 mins.	11 mins.
402 - 500K	506	190 mins.	13 mins.

S/370 Performance in August

There were twenty-one days with no interruption of service in August. The other ten days averaged almost two a day; nine software, eight hardware, and two electrical failures. All in all service was much better than July, but with plenty of room for improvement.

	July 1975	August 1975
Hours available (Batch and TSO)	385.3	408.4
Daily average of hours available	12.8	13.2
Hours of unattended TSO available	253.7	273.7
Daily average of hours of unattended TSO	8.2	8.8
Hours lost due to failures	27.3	8.6
Hardware	22.9	5.2
Software	4.4	3.4
Number of failures	42	19
Mean time between failures	15.2 hours	35.8 hours
Mean time to repair	39 minutes	27 minutes

TSO Statistics for August

Number of LOGON sessions	7,821	
Total CONNECT time	5,083	hours
Average CONNECT time per session	39	minutes
Total CPU time	20	hours
Average CPU time per session	0.15	minutes
Total disk I/O operations	9.3	million
Average disk I/O operations per session	1,189	
Total teleprocessing I/O operations	2.3	million
Average teleprocessing I/O operations per session	304	
Total K byte minutes of interactive memory use	.0.71	million
Average K byte minutes per session	90	

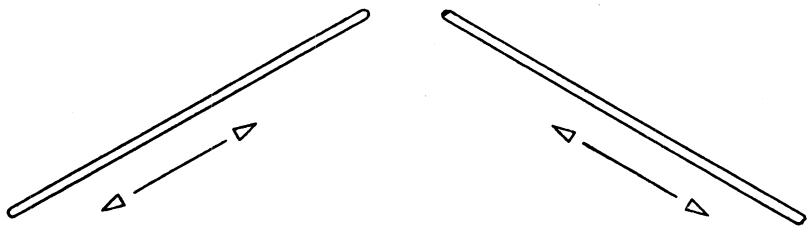

JOB PROCESSING SYSTEM STATISTICS FOR AUGUST

The statistics reported include only M.I.T.'s use of the facility. Each of the averages is based on the actual number of jobs contributing to the statistic.

Number of M.I.T. jobs processed on the 168	19,889	
Number of M.I.T. job steps processed on the 168	42,203	
Total CPU time used	78.0	hours
Average CPU time used per job	0.26	minutes
*Total run time used (estimated)	236	hours
Average run time per job	0.71	minutes
Total number of I/O operations by type		
CTC	9.0	million
Tape	2.1	million
Disk	16.0	million
Average number of I/O operations per job		
CTC	518	
Tape	1,073	
Disk	896	
Total memory minutes used	3.3	million K bytes
Average memory minutes used per job	186	K bytes
Total number of lines printed	26.4	million
Average number of lines printed per job	1,349	
Total number of cards read	7.0	million
Average number of cards read per job	352	
Total number of cards punched	0.62	million
Average number of cards punched per job	437	
Total number of Remote Job I/O units	108	
Average number of RJ I/O units per job	108	
Total number of setup jobs	4,306	
Total number of setup tapes	2,835	
Total number of setup disks	2,653	

*Run time is roughly equivalent to the real time the jobs would require if run alone under a PCP or MFT system.

TURN-AROUND STATISTICS FOR AUGUST

The following table reports turn-around figures for the Job Processing System for the month of August, where turn-around is measured as the interval between the time the job enters the system via the card reader and the time it is purged. The statistics are for M.I.T. jobs only.

Turn-around by SRI

		LOW	STANDARD	HIGH	ALL*
Number of jobs		6,383	11,906	936	19,890
Average Turn-around		91 mins.	30 mins.	21 mins.	116 mins.
Percentage of Jobs Returned	25%	11 mins.	6 mins.	5 mins.	7 mins.
	50%	42 mins.	13 mins.	10 mins.	23 mins.
	75%	120 mins.	29 mins.	25 mins.	58 mins.
	95%	359 mins.	104 mins.	59 mins.	359 mins.

* Including weekly

Setup vs. Non-setup Jobs

The average turn-around for 15,583 non-setup jobs was . . . . . 99 mins.
The average turn-around for 4,307 setup jobs was . . . . . .176 mins.

Turn-around by Region

Region	Number of Jobs	Average Turn-around	Median Turn-around
0 - 100K	2,280	124 mins.	14 mins.
102 - 200K	9,816	88 mins.	14 mins.
202 - 300K	4,776	128 mins.	15 mins.
302 - 400K	1,730	146 mins.	25 mins.
402 - 500K	861	172 mins.	37 mins.