

EXPERIMENTS WITH THE
BURROUGHS B 3500 COMPUTER SYSTEM
USING A SYNTHETIC WORKLOAD

K. Sreenivasan
J. E. Esposito

JANUARY 1975

Prepared for

DEPUTY FOR COMMAND AND MANAGEMENT SYSTEMS

ELECTRONIC SYSTEMS DIVISION
AIR FORCE SYSTEMS COMMAND
UNITED STATES AIR FORCE
L. G. Hanscom Field, Bedford, Massachusetts



Approved for public release;
distribution unlimited.

Project No. 572S
Prepared by
THE MITRE CORPORATION
Bedford, Massachusetts
Contract No. F19628-73-C-0001

BEST AVAILABLE COPY

ADH004804

When U.S. Government drawings, specifications, or other data are used for any purpose other than a definitely related government procurement operation, the government thereby incurs no responsibility nor any obligation whatsoever; and the fact that the government may have formulated, furnished, or in any way supplied the said drawings, specifications, or other data is not to be regarded by implication or otherwise, as in any manner licensing the holder or any other person or corporation, or conveying any rights or permission to manufacture, use, or sell any patented invention that may in any way be related thereto.

Do not return this copy. Retain or destroy.

REVIEW AND APPROVAL

This technical report has been reviewed and is approved for publication.




THOMAS V. REILLY, TSgt, USAF
Field Support Division



WILLIAM J. LETENDRE
Field Support Division

FOR THE COMMANDER



ROBERT W. O'KEEFE, Colonel, USAF
Director, Information Systems
Technology Applications Office
Deputy for Command and Management Systems

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ESD-TR-74-312	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) EXPERIMENTS WITH THE BURROUGHS B 3500 COMPUTER SYSTEM USING A SYNTHETIC WORKLOAD		5. TYPE OF REPORT & PERIOD COVERED
7. AUTHOR(s) K. Sreenivasan J. E. Esposito		6. PERFORMING ORG. REPORT NUMBER MTR-2859
9. PERFORMING ORGANIZATION NAME AND ADDRESS The MITRE Corporation Box 208 Bedford, Mass. 01730		8. CONTRACT OR GRANT NUMBER(s) F19628-73-C-0001
11. CONTROLLING OFFICE NAME AND ADDRESS Deputy for Command and Management Systems Electronic Systems Division, A. F. S. C. L. G. Hanscom Field, Bedford, Mass. 01730		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Project No. 572S
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE JANUARY 1975
		13. NUMBER OF PAGES 48
		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) ACCOUNTING DATA COMPUTER PERFORMANCE HARDWARE MONITOR SYNTHETIC WORK LOAD		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) An application of synthetic workloads in the evaluation of the performance of batch processing, multiprogrammed, real-memory computer systems is described. Results of controlled experiments with a Burroughs B 3500 computer system in a dedicated environment using a test workload are presented. The test workload is constructed by matching the joint probability distribution of the selected real workload characteristics. They are the core size of the program, the number of files, I/O activity to disk and the CPU time. These characteristics are determined from the system accounting package, LOGGER. The		

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

20. (Cont.)

accounting data, collected by LOGGER during the controlled experiments with the synthetic workload, are used to determine the utilization of the processor and the I/O channel. These utilization values compare favorably with those measured by a DYNAPROBE 7900 hardware monitor.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

TABLE OF CONTENTS

	<u>Page</u>
LIST OF ILLUSTRATIONS	2
LIST OF TABLES	2
SECTION I INTRODUCTION	4
SECTION II CONSTRUCTION OF TEST WORKLOAD	6
GENERAL DESCRIPTION	6
PREPROCESSING OF SYSTEM LOG FILE DATA	8
LOGGER Data	8
Data Structure	8
ACCESSABILITY	9
COLLECTION OF JOB-ORIENTED DATA	12
Detailed Format	12
Method	12
WORKLOAD CHARACTERISTICS	12
TEST WORKLOAD CHARACTERISTICS	15
SYNTHETIC WORKLOAD	21
CALIBRATION AND INVERSION	21
CLOSURE	25
SECTION III EXPERIMENTS WITH THE TEST WORKLOAD	29
PURPOSE	29
EXPERIMENTAL RESULTS	29
HARDWARE MONITORING	34
SECTION IV SUMMARY	38
APPENDIX I CALCULATION OF THE RESOURCE UTILIZATION	40
REFERENCES	46

LIST OF ILLUSTRATIONS

<u>Figure Number</u>		<u>Page</u>
1	File Close and Open Records	10
2	End of Job and Beginning of Job Records	11
3	Job Summary Record Used in Workload Characterization	13
4	B 3500 Hardware Configuration	31
5	Variation of the Degree of Multiprogramming During Experiment 4	33

LIST OF TABLES

<u>Table Number</u>		<u>Page</u>
I	Typical B 3500 Job Summary Records (Fixed portion only)	14
II	Histogram of Processor Time	16
III	Characteristics of the Test Workload	20
IV	Joint Probability Distribution of the Real Workload	22
V	Experimental Results Used in Calculating the Synthetic Program	26
VI	Overall Summary of the Experimental Results	35
VII	DYNAPROBE Hardware Monitor Reduction Output	37

LIST OF TABLES (Continued)

<u>Table Number</u>		<u>Page</u>
VIII	Comparison of the Processor Utilization Values	41
IX	Comparison of the Channel Utilization Values	45

SECTION I

INTRODUCTION

The performance of computer systems is usually evaluated for the purposes of determining its present value, methods of improving it and predicting the effects of changes in either the workload or the system. One method of studying these factors is to conduct experiments with an existing system using a test workload. This report describes the results of and experiences from an experimental investigation with a Burroughs B 3500 computer system at ESD, Hanscom Air Force Base.

In order to conduct these experiments there is a need for a stable, reproducible and flexible test workload that imitates the real workload with reasonable fidelity but in an abbreviated form. It is necessary that the test workload be stable and reproducible so that the experimental results are interpreted correctly; the test workload be flexible so that the characteristics of the test workload may be varied; and the test workload be a representative model of the real workload so that valid conclusions may be drawn from its use. All these requirements can be fulfilled by using synthetic programs in the test workload.

The test workload used in the experiments is constructed by matching the joint probability density for the selected workload characteristics of the real workload with those of the test workload. The selected characteristics are the processor time, the I/O activities to disk, the number of files and the amount of core used by each job. These characteristics are determined from the LOGGER accounting data maintained by the Master Control Program (MCP-V) of the Burroughs B 3500 system. The details of the various types of records collected by LOGGER are studied and special conversion

packages are implemented to create a job summary record for every job processed. These summary records are used to determine the test workload characteristics.

In these experiments the utilization of the processor and the channel is measured using a DYNAPROBE 7900* hardware monitor. These utilization values are also calculated using the LOGGER data. The two sets of values compare favorably. This is a significant finding of this investigation. The utilization values of the system resources are essential in the evaluation of the computer performance. This study indicates that the accounting data (LOGGER) provides a ready and economical source for determining the resource utilization. A system manager can conveniently assess the system performance by periodic processing of the accounting data to determine the resource utilization.

The test workload is only a model of the real workload and is constructed by striking a compromise between the needs of representativeness and physical constraints such as the processor time used. Such a compromise necessarily renders the synthetic workload not completely representative of all the features of the real workload. This is indeed the case in our study where the test workload is constructed using only four of the many characteristics recorded by the accounting data. The tape I/O activities are not represented because it was decided not to consider the human factors involved in the tape handling. Special measurements have to be made to include the tape I/O activities in the test workload. These measurements relate to the method of handling tapes.

The construction of the test workload is outlined in Section II. Section III describes the experiments. The report is summarized in Section IV and the method of calculating the resource utilization is described in the Appendix.

*Comress, Incorporated.

SECTION II

CONSTRUCTION OF TEST WORKLOAD

GENERAL DESCRIPTION

A workload is defined as the collection of all the individual jobs that are processed by a computer system during a specified period of time. The computer system can be considered as a collection of resources upon which the workload places certain demands. The magnitude of these demands can be viewed as the characteristic variables of the real workload. A job, such as a compilation, matrix inversion, or sort-merge can be described by a set of these variables whose magnitudes will vary from one type of job to another, and from one computer system to another.

In this method of characterizing the workload it is assumed that since the system does not recognize the type of job, two jobs reflecting the same value for the characteristic variables are treated identically by the operating system. This assumption is reasonable since the operating system classifies jobs on the basis of the demands they place on the system resources.

Many computer installations maintain a system accounting package that, for charging purposes, gathers information about the use of various system resources. This information provides a ready source of data from which many of the workload characteristics can be derived. To determine the characteristics not available from the accounting data, hardware and software monitors must be used.

In many installations a significant part of the I/O activities may be directed to tape. This was the case in this study where the workload of the B 3500 (ESD, Hanscom Air Force Base) was studied. In such cases the accounting data seldom record the details of the I/O activities to tape. Examples of these details are: number of tape mounts

and dismounts, time per mount, number of jobs requesting the same tape drive, and forced idle time for a job because of conflicts in tape drive availability. These details are significant and contribute to the overall performance of the computer system. Although these details are system dependent they can best be described as human factors that influence the throughput of the workload. Mounting and dismounting of tapes are not only influenced by the number of tapes and the location of the tape library but also by the number of operators available.

The method of characterizing the workload described in this report does not consider the human factors. This study used the total number of I/O activities to tape or to disk as a measure of the I/O load. This quantity, admittedly, neglects the time for mount and dismount and this is a limitation in this method of characterization.

The tape I/O can be accounted for by creating tape files in the synthetic program. Initial calibration experiments must be conducted to relate the tape block counts (the number of blocks read or written to tape) with the corresponding synthetic program parameters. In the case of tape files, not all the block counts (recorded by the accounting package) lead to data transfer; some of them result in tape positioning, label checking and tape rewinding. The fraction of the total block count resulting in data transfer can be determined in the preliminary calibration experiments. The tape handling in the real workload should then be studied to estimate the time taken for responding to the mount request and the number of tape files processed by individual jobs. Based on these studies it is possible to execute the synthetic program creating tape files and introduce known amounts of delay to simulate tape handling.

PREPROCESSING OF SYSTEM LOG FILE DATA

LOGGER Data

The LOGGER provided the source of performance data for this report. LOGGER is the accounting package provided as a part of the MCP operating system for the B 3500. LOGGER records the following event-oriented statistics. They are usually referred to as the Type X records where X is any one of the following.

Type 0	File Close Record
Type 1	File Open Record
Type 2	End of Job Record
Type 3	Long Schedule Record
Type 4	Short Schedule Record
Type 5	Comment Record
Type 6	Beginning of Job Record
Type 7	Idle Time Record
Type 8	Halt/load Record
Type 9	Filler Record

Data Structure

The accounting package collects and records data in the event-oriented format as and when the events take place. Typically they are: job A begins; job A opens file x; job B begins; job A closes file x; etc. They are also referred to as the raw data. These data are uniquely identified by the job-log number (or the job ID) and using this field as the key the raw data are summarized into job-oriented format, also referred to as the summary data. The job-oriented data, essentially, are the magnitudes of the demands placed on the various system resources by the individual jobs. It should be noted that one job-oriented record is obtained by summarizing several event-oriented records.

The performance variables of interest for this workload characterization were taken from four of the LOGGER record types: Type 0, File Close Record; Type 1, File Open Record; Type 2, End of Job Record; and Type 6, Beginning of Job Record. The record layouts for each of these types are shown in Figures 1 and 2.^[1]

The standard internal character set for the B 2500/B 3500 is the 8-bit Extended Binary Coded Decimal Interchange Code (EBCDIC). Internal storage for the accounting data is organized into 4-bit digits and is processed under two formats: unsigned 4-bit numeric and 8-bit alphanumeric. Most of the data uses the unsigned 4-bit numeric format which is the Burroughs structure for high density storage of data.^[2] The voluminous quantity of event-oriented records on system activity must be summarized into usable job-oriented summary records. An existing program written in PL/I for the IBM 370/155 (MITRE, Bedford) was modified and used to create job summary records. It was necessary, therefore, to convert the Burrough's LOGGER data into IBM compatible data types. There is no conversion necessary for the 8-bit alphanumeric data but the unsigned 4-bit numeric must be changed by adding a sign and aligning digits on 8-bit (character) boundaries.

ACCESSIBILITY

The event-oriented data collected by the LOGGER is readily accessible. The data set receiving this raw data is normally a disk file. Log procedures of the B 3500 accounting system store information into one of three disk data sets called System Log Files. The raw data of interest in this analysis is in the system run (current) log file and is named RLOG. The disk files are periodically emptied onto tape volumes for storage. The RLOG data, consequently, is available from either the active disk file or the stored tape file.

CLOSE		OPEN	
<u>Position</u>	<u>Contents</u>	<u>Position</u>	<u>Contents</u>
0	Type Code	0	Type Code
1-4	Reserved	1-4	Reserved
5-8	Log ID Number	5-8	Log ID Number
9-14	Date (MMDDYY)	9-14	Date (MMDDYY)
15-22	Time (msecs)	15-22	Time (msecs)
23	Subtype	23	Subtype
24-35	File ID	24-35	File ID
36-47	Multi-File ID	36-47	Multi-File ID
48-49	File Number	48-49	File Number
50-51	Primary I/O Channel	50-51	Primary I/O Channel
52	Unit Number	52	Unit Number
53-54	Hardware Type Used	53-54	Hardware Type Requested
55	Supplementary Hardware Code	55	Supplementary Hardware Code
56-58	Reel Number	56-58	Reel Number
59-63	Physical Tape Number	59-63	Creation Date
64-65	Close Type Code	64-65	Cycle Number
66-73	Logical Record Count	66-70	Maximum Record Length
74-81	Physical Record Count	71-73	Records per Block
82-84	Exception Count	74-79	Maximum Block Size
85-86	Number of Disk Areas Used	80	Buffer Access Technique
87-94	Disk End-of-File Pointer	81	File Label Convention
95	Memory for Disk File Headers	82	Number of Alternate Areas
96-99	Reserved	83	OPEN Type Code
		84	Recording Mode
		85	Blocking Technique
		86	Special Forms Indicator
		87-89	Save Factor
		90-96	Disk Segments per Area
		97	Disk Access Technique
		98	Disk File Header Block Count
		99	Reserved

Figure 1. File Close and Open Records

END		BEGIN	
<u>Position</u>	<u>Contents</u>	<u>Position</u>	<u>Contents</u>
0	Type Code	0	Type Code
1-4	Reserved	1-4	Reserved
5-8	Log ID Number	5-8	Log ID Number
9-14	Data (MMDDYY)	9-14	Date (MMDDYY)
15-22	Time (msecs)	15-22	Time (msecs)
23	Subtype	23	Subtype
24-35	Program ID	24-35	Program ID
36-47	Multi-program ID	36-47	Multi-program ID
48-49	Job Number	48-49	Job Number
50-51	Primary I/O Channel	50-55	Disk Segments in Program
52	Unit Number	56-61	User Charge Number
53-54	Hardware Type	62-64	Core Required
55	Supplementary Hardware Type	65-66	Number of Files
56-61	Overlay Count	67-68	Number of Disk Files
62-63	Finish Code	69	Execution Code
64-73	Reserved	70	Reserved
74-81	Direct Processor	71	Supplementary Execution Code
82-89	Prorated Processor Time	72	MCP Intrinsic Flag
90-97	Accumulated Program I/O Wait Time	73	Reserved
98-99	Reserved	74-79	Date Compiled
		80-99	Reserved

Figure 2. End of Job and Beginning of Job Records

COLLECTION OF JOB-ORIENTED DATA

Detailed Format

The job summary record in Figure 3 consists of performance variables of interest for a workload characterization. A job record to be used for another purpose, e.g., billing reports, would probably embody a significantly different set of information. The Log ID Number, Job Start Time, and Core Required fields are taken from BOJ record type 6, Figure 2. EOJ record type 2, Figure 2, yields Job End Time, Direct Processor Time, Prorated Processor Time, Hardware Type, and job Log ID. Elapsed Time is computed using EOJ and BOJ times taken from type 2 and type 6 records respectively. The remaining fields itemizing files, devices, channels, and block counts are taken from the file close records, type 0, Figure 1.

Method

Each event-oriented record is read in the order it was created and checked for type 6, the BOJ record. A new BOJ record defines a unique job identifier by means of the Log ID Number. All subsequent records pertaining to a particular job are found by using this identifier. As each job starts it is added to the mix and with each EOJ the job is removed by updating the mix count. All file open and close records for a particular job are processed by counting the number of different files, devices, and channels that appear and by breaking down the block count into file, device, and channel sub-totals. The EOJ record finishes the summary cycle for a job. The processor time is extracted from this final record and the completed job summary record of Figure 3 is recorded. Table I shows a partial list of summary records created by processing one month of raw accounting data.

WORKLOAD CHARACTERISTICS

The system dependent characteristics of a month's workload processed by B 3500 were analyzed in order to determine the test workload

<u>Position</u>	<u>Length (Bytes)</u>	<u>Data Type</u>	<u>Contents</u>
0-2	3	PD	Log ID Number
3-7	5	PD	Job Start (msecs)
8-12	5	PD	Job End Time (msecs)
13-17	5	PD	Elapsed Time (msecs)
18-22	5	PD	Direct Processor Time (msec)
23-27	5	PD	Prorated Processor Time (msec)
28-32	5	PD	Total Processor Time (msec)
33-34	2	PD	Core Required (kilo-bytes)
35-39	5	PD	Total Block Count
40-41	2	B	Number of Files (L)
42-43	2	B	Number of Devices (M)
44-45	2	B	Number of Channels (N)
46-V	L x 6	CH	File Names
V	M x 2	PD	Device Types
V	N x 2	PD	Channels
V	M x 2	B	Files per Device
V	N x 2	B	Files per Channel
V	N x 2	B	Devices per Channel
V	L x 5	PD	Blocks per File
V	M x 5	PD	Blocks per Device
V	N x 5	PD	Blocks per Channel

Key: PD - Packed Decimal
 B - Binary
 CH - Character
 L - Number of Files
 M - Number of Devices
 N - Number of Channels
 V - Variable

Figure 3. Job Summary Record Used in Workload Characterization

LOG ID NO	START TIME	END TIME	ELAPSED TIME	DIRECT TIME	PRORATED TIME	BLK COUNT	CUKE REQ (K BYTES)	NU OF FILES	NU OF DEVICES	NO OF CHANNELS
60	33604.523	33607.402	2.879	0.290	0.276	2	2	1	1	1
61	33345.133	33744.479	399.346	12.528	51.103	95	29	2	2	1
62	33454.392	33752.466	298.074	33.539	22.142	268	29	2	2	1
63	33618.921	33759.455	270.534	9.515	24.390	883	4	5	2	2
64	33704.080	34208.336	504.256	144.085	36.743	6235	23	10	2	2
65	33319.293	34293.610	474.317	52.349	35.741	6197	4	10	2	2
66	34364.520	34705.627	341.107	42.677	18.519	4961	4	6	2	2
67	34316.840	34823.367	516.541	211.895	52.617	8972	23	16	2	2
68	34025.088	34962.730	137.644	4.195	15.699	371	4	3	2	2
69	35008.039	35407.028	519.712	254.571	43.563	8926	23	10	2	2
70	35522.602	36080.003	207.416	3.247	10.279	2	18	2	2	2
71	35047.328	36310.312	525.345	60.484	51.321	3975	23	16	2	2
72	36324.461	36423.194	468.492	83.553	46.669	1322	29	3	3	2
73	36554.199	36615.953	98.753	13.813	13.636	1215	4	2	2	2
74	36515.623	36705.637	61.794	3.100	3.843	271	3	2	2	2
75	37185.357	37313.372	87.015	3.145	13.041	284	4	3	2	2
76	37312.496	37399.301	87.015	12.802	3.470	407	19	9	3	2
77	37562.728	37653.775	3.047	16.480	6.000	444	11	2	2	1
78	37671.607	37797.350	33.125	5.581	1.079	2	4	1	1	1
79	37930.556	38016.245	85.678	21.619	5.252	2512	4	4	2	2
80	38135.555	38183.709	48.154	3.275	5.360	2	10	2	2	2
81	38196.704	38254.673	58.969	35.359	17.221	1321	13	9	4	3
82	38473.748	38531.237	57.489	17.430	13.224	1361	4	4	2	2
83	38672.492	38800.623	7.225	13.634	14.407	220	13	7	3	3
84	38942.513	39010.257	15.744	1.234	6.225	99	4	2	2	2
85	39155.474	39250.312	94.839	3.514	6.257	181	13	9	2	2
86	39354.644	39415.227	21.581	1.200	1.200	93	5	2	2	2
87	39544.598	39515.227	2301.529	1094.657	149.270	34004	13	13	2	2
88	39144.538	39650.739	552.151	203.609	7.323	12998	13	47	4	3
89	39155.474	39750.691	595.217	205.803	6.461	6730	24	18	4	2
90	39659.309	39950.312	117.003	0.721	1.042	19	4	3	2	2
91	39970.233	39982.644	11.405	1.666	6.420	42	20	3	1	1
92	40142.474	40250.971	6.477	7.521	3.304	1	3	2	2	2
93	40115.681	40230.402	122.721	7.350	9.319	339	3	2	2	2
94	39715.837	40274.246	558.409	103.269	23.961	747	29	5	3	1
95	39779.408	40325.611	44.13	3.166	3.230	38	20	3	1	1
96	40225.574	40229.400	3.826	0.160	1.162	7	4	1	1	1
97	40391.277	40430.172	38.895	3.059	1.162	49	4	3	2	2
98	40406.513	40417.215	10.762	1.739	1.645	42	20	3	1	1
99	40417.739	40425.664	7.925	0.737	6.542	22	19	4	3	2
100	40479.430	40577.402	25.981	2.849	1.703	51	15	3	2	2
101	40577.402	40584.089	11.680	1.655	6.780	122	2	2	2	2
102	40592.640	40603.344	11.298	1.624	6.037	394	4	3	2	2
103	40607.801	40673.751	65.950	4.736	5.239	11	4	3	2	2
104	40549.700	40693.619	140.321	3.435	11.431	394	15	3	2	2
105	40550.252	40701.218	150.966	1.623	12.250	6	13	3	2	2
106	40550.591	40701.412	150.821	2.125	1.639	51	15	3	2	2
107	40749.234	40759.144	19.910	2.100	1.639	11	4	3	2	2
108	40759.221	40802.564	43.343	3.441	1.637	383	4	3	2	2
109	40759.221	40921.901	23.043	3.270	2.053	383	5	3	2	2
110	40979.010	41001.122	102.100	11.014	4.704	894	14	3	3	2
111	40702.602	41118.264	355.562	33.252	2.704	2341	29	13	3	3
112	41197.195	41203.353	6.158	1.602	6.220	216	3	2	2	2
113	41300.215	41300.215	45.412	7.491	2.900	450	13	7	3	2
114	41194.273	41520.041	431.765	11.207	35.410	1160	4	9	2	2

Table I --- Typical B 3500 Job Summary Records (Fixed portion only)

characteristics. Single variable histograms were constructed to isolate the major variables. The following histograms were studied:

- (1) Direct Time
- (2) I/O activities to disk
- (3) I/O activities to tape
- (4) Core used
- (5) Number of files
- (6) Number of channels

As an example, the histogram of the direct time is tabulated in Table II. Based on these histograms it was decided to characterize a job by the following four variables.

- (1) Direct Time, seconds
- (2) Number of block counts to disk
- (3) Core used
- (4) Number of files

The direct time is the demand placed on the CPU and is the processor time in the user state or the normal state (or the problem program state). The number of block counts is the number of blocks of data exchanged between main storage and the disk. The other two variables are self-explanatory. These four job-oriented characteristics were determined from the derived LOGGER summary data. The general method of the statistical analyses to determine the workload characteristics appears in Reference 3.

TEST WORKLOAD CHARACTERISTICS

The test workload characteristics were determined by matching the four-dimensional probability density of the test workload with that of the real workload. The details of this method appear in Reference 4 and will be briefly presented here.

SUMMARY OF CPU(DIRECT) TIME IN SECS.

XMIN= 0.0000 XMAX= 200.0000
 ANIN= 0.0000 AMAX= 4305.7000
 MEAN VALUE= 56.0307
 STD DEVIATION= 152.0641
 MAXIMUM FREQUENCY AT X= 0.0000 WITH 2011 JOBS

HISTOGRAM OF CPU(DIRECT) TIME IN SECS.

INTERVAL	COUNT	DISTRIBUTION	DENSITY
0.0000	2011	0.000000	0.000000
2.0000	710	0.001293	0.000000
4.0000	455	0.002502	0.000000
6.0000	327	0.004200	0.000000
8.0000	262	0.003091	0.000000
10.0000	199	0.002509	0.000000
12.0000	176	0.002029	0.000000
14.0000	169	0.001591	0.000000
16.0000	161	0.001091	0.000000
18.0000	143	0.000807	0.000000
20.0000	122	0.000566	0.000000
22.0000	95	0.000207	0.000000
24.0000	89	0.001144	0.000000
26.0000	84	0.001078	0.000000
28.0000	79	0.001144	0.000000
30.0000	74	0.001078	0.000000
32.0000	87	0.001115	0.000000
34.0000	80	0.000204	0.000000
36.0000	57	0.000732	0.000000
38.0000	83	0.000800	0.000000
40.0000	87	0.000800	0.000000
42.0000	103	0.001245	0.000000
44.0000	53	0.000805	0.000000
46.0000	49	0.000207	0.000000
48.0000	35	0.000400	0.000000
50.0000	43	0.000504	0.000000
52.0000	37	0.000475	0.000000
54.0000	44	0.000500	0.000000
56.0000	36	0.000600	0.000000
58.0000	30	0.000404	0.000000
60.0000	44	0.000500	0.000000
62.0000	45	0.000500	0.000000
64.0000	36	0.000404	0.000000
66.0000	34	0.000304	0.000000
68.0000	42	0.000500	0.000000
70.0000	21	0.000200	0.000000
72.0000	27	0.000304	0.000000
74.0000	41	0.000500	0.000000
76.0000	34	0.000304	0.000000
78.0000	30	0.000304	0.000000
80.0000	34	0.000304	0.000000

Table II -- Histogram of Processor Time

82.0000	36	0.004623	0.470647	***
84.0000	31	0.003963	0.395647	***
86.0000	37	0.004745	0.460772	***
88.0000	27	0.003404	0.340271	***
90.0000	21	0.002860	0.282474	***
92.0000	31	0.003973	0.346524	***
94.0000	26	0.003342	0.282474	***
96.0000	13	0.001603	0.251474	***
98.0000	15	0.002244	0.252474	***
100.0000	22	0.002763	0.252474	***
102.0000	28	0.003280	0.261474	***
104.0000	24	0.002864	0.254711	***
106.0000	16	0.002045	0.258736	***
108.0000	13	0.001603	0.254711	***
110.0000	16	0.002045	0.254711	***
112.0000	13	0.001603	0.254711	***
114.0000	11	0.001244	0.254711	***
116.0000	17	0.002174	0.254711	***
118.0000	17	0.002174	0.254711	***
120.0000	14	0.001791	0.254711	***
122.0000	14	0.001791	0.254711	***
124.0000	22	0.002823	0.254711	***
126.0000	20	0.002504	0.254711	***
128.0000	19	0.002441	0.254711	***
130.0000	11	0.001474	0.254711	***
132.0000	15	0.002374	0.254711	***
134.0000	20	0.002503	0.254711	***
136.0000	9	0.001144	0.254711	***
138.0000	13	0.001603	0.254711	***
140.0000	21	0.002506	0.254711	***
142.0000	17	0.002122	0.254711	***
144.0000	14	0.001765	0.254711	***
146.0000	22	0.002823	0.254711	***
148.0000	11	0.001474	0.254711	***
150.0000	10	0.001244	0.254711	***
152.0000	6	0.000763	0.254711	***
154.0000	7	0.000845	0.254711	***
156.0000	14	0.001791	0.254711	***
158.0000	9	0.001144	0.254711	***
160.0000	13	0.001603	0.254711	***
162.0000	15	0.001822	0.254711	***
164.0000	13	0.001603	0.254711	***
166.0000	9	0.001144	0.254711	***
168.0000	2	0.000244	0.254711	***
170.0000	6	0.000763	0.254711	***
172.0000	8	0.001244	0.254711	***
174.0000	9	0.001244	0.254711	***
176.0000	7	0.000845	0.254711	***
178.0000	7	0.000845	0.254711	***
180.0000	2	0.000244	0.254711	***
182.0000	9	0.001144	0.254711	***
184.0000	7	0.000845	0.254711	***
186.0000	6	0.000763	0.254711	***
188.0000	7	0.000845	0.254711	***
190.0000	6	0.000763	0.254711	***
192.0000	4	0.000544	0.254711	***
194.0000	1	0.000122	0.254711	***
196.0000	3	0.000371	0.254711	***
198.0000	7	0.000845	0.254711	***
200.0000	512	0.004745	0.470647	***

Table II -- Histogram of Processor Time (Continued)

The B 3500 workload is characterized by the following four variables.

- X_1 - Number of files
- X_2 - Core used
- X_3 - Direct time
- X_4 - Number of block counts to disk

A job, therefore, may be regarded as a point in the multi-dimensional space, with the co-ordinates representing each of the four demands. The workload then becomes an ensemble of points in this co-ordinate system.

In general, if X_1, X_2, X_3 , and X_4 are the four variables used to describe the workload and N_{ijkl} is the number of jobs in the cell* with co-ordinates $X_1 = x_1^{(i)}$, $X_2 = x_2^{(j)}$, $X_3 = x_3^{(k)}$, $X_4 = x_4^{(l)}$ then the workload can be described as a multi-variate probability density function and its value is given by:

$$p_{ijkl} = \frac{N_{ijkl}}{N_{tot}} \quad i, j, k, l = 1, 2, 3 \dots L \quad (2.1)$$

where p_{ijkl} is the probability of finding a job in the (i-j-k-l)-th cell and N_{tot} is the total number of jobs in the workload. L is the number of cells along each co-ordinate axis.

The test workload is constructed by matching the joint probability density of the four variables for the test workload with that of the

* The magnitude of the variable corresponding to each cell is taken to be its value at the mid-point of the interval.

real workload. Denoting the test workload characteristics by primed quantities, we have

$$N'_{ijkl} = p_{ijkl} N'_{tot} \quad (2.2)$$

This equation can be physically interpreted as matching the joint probability density of the test workload with that of the real workload. Constraints on the total number of jobs in the test workload or on the total CPU time for processing it, can be imposed in the determination of N'_{ijkl} . In this study the constraint was placed on the total CPU time, namely, 2000 seconds. The characteristics of the test workload are derived using equation (2.2) and they are tabulated in Table III. The test workload consists of 87 jobs.

The total number of cells used in the determination of p_{ijkl} and N'_{ijkl} is governed by the value of L , or the number of cells along each co-ordinate axis. In our study the value of $N_{tot} = 7776$. A very large value of L will result in a large value of the total number of cells (L^4). As seen from equation (2.1) there are L^4 values of p_{ijkl} (for example, p_{1111} , p_{1112} , ..., p_{2222} , ..., p_{LLLL}). Since the total number of jobs (N_{tot}) in the ensemble is fixed, large values of L result in small values of p_{ijkl} . It should be pointed out that there is a rounding-off performed in determining the N'_{ijkl} from equation (2.2). If the value of N'_{ijkl} is less than 0.5 it will be rounded-off to 0. A large value of L will tend to diffuse the distribution and will result in a multi-dimensional picture that is not useful in describing the workload. Before selecting L the single variable histograms should be studied to determine the values of the variables that correspond to the most frequently occurring jobs. It is also not necessary to choose the same value of L along all the co-ordinate axes. The number and location of the cells is influenced by the single variable histograms. Based on these considerations the ensemble in

Table III
Characteristics of the Test Workload

Group	Number of Jobs (N_{ijkl})	Core (K-bytes)	Number of Files	CPU (sec.)	Block Count
1	14	4.5	3	2	25
2	4	4.5	3	4	200
3	5	4.5	3	5	200
4	2	4.5	3	9	1000
5	2	4.5	3	29	1000
6	1	4.5	3	73	5000
7	1	4.5	6	6	200
8	1	4.5	6	9	1000
9	1	4.5	6	31	1000
10	1	4.5	6	43	5000
11	1	6.5	11	33	1000
12	1	6.5	11	81	5000
13	5	17.5	3	3	25
14	3	17.5	3	6	200
15	1	17.5	3	9	1000
16	3	17.5	3	34	1000
17	1	17.5	3	78	5000
18	1	17.5	6	1	25
19	2	17.5	6	6	25
20	4	17.5	6	7	200
21	1	17.5	6	11	1000
22	5	17.5	6	31	1000
23	2	17.5	6	77	5000
24	2	18.0	11	7	200
25	6	18.0	11	34	1000
26	1	18.0	11	41	5000
27	1	18.0	11	75	1000
28	6	18.0	11	74	5000
29	1	29.5	3	4	25
30	2	29.5	3	30	200
31	1	29.5	6	8	200
32	1	29.5	6	30	1000
33	1	29.5	6	81	5000
34	1	30.0	11	29	1000
35	2	30.0	11	72	5000

this study was divided into 144 cells by choosing 3 values of the number of files, 3 values for the core size, 4 values for the direct time, and 4 values for the block count to disk. Table IV tabulates the 144 values for p_{ijkl} .

SYNTHETIC WORKLOAD

The eighty-seven jobs in the test workload are realized by using a synthetic program. A COBOL program was designed and implemented for this purpose. It is not practical to use one synthetic program that covers a large range of variation in core size and number of files; further such a synthetic program will not be flexible. It is simple and easy to implement several synthetic programs to represent the core size and the number of files which are specified as compile time parameters. The four variables to be represented are divided into two groups. They are:

Run time variables - Direct time and number of block counts to disk.

Compile time variables - Core size and number of files.

Two parameters were built into the synthetic program for varying the direct time and the number of block counts to disk. There are nine combinations of core size and the number of files and nine separate COBOL programs were compiled to reflect these nine combinations. The nine programs were stored in auxiliary storage, retrieved at run time and executed using the run time parameters.

CALIBRATION AND INVERSION

The nine values of the core size and number of files combinations were made equal to the corresponding values of the test workload. This eliminated the need for calibrating the synthetic program for these two parameters. Each one of the nine programs has two parameters for varying the direct time and the number of block counts

I - CORE
J - FILES
K - CPU
L - IO

	* L	25	200	1000	5000
	*				
	K	*****	*****	*****	*****
	*				
I = 4	1000 *	1056	282	0	0
J = 3	6000 *	29	372	189	0
	30000 *	4	10	164	56
	72000 *	2	0	30	102

	* L	25	200	1000	5000
	*				
	K	*****	*****	*****	*****
	*				
I = 4	1000 *	23	7	0	0
J = 6	6000 *	1	100	97	0
	30000 *	0	1	103	67
	72000 *	0	0	0	56

	* L	25	200	1000	5000
	*				
	K	*****	*****	*****	*****
	*				
I = 4	1000 *	1	0	0	0
J = 11	6000 *	0	21	19	0
	30000 *	0	1	85	53
	72000 *	0	0	0	92

Table IV -- Joint Probability Distribution of the Real Workload. For example, the pair of numbers (1056, 0.136) for I = 4, J = 3, K = 1000, L = 25 represent N_{ijkl} and P_{ijkl} respectively.

	* L	25	200	1000	5000
I = 17 J = 3	*				
K	*****	*****	*****	*****	*****
1000 *	412 .053	35 .C04	0 .000	0 .000	
6000 *	17 .002	198 .C25	73 .009	0 .000	
30000 *	3 .000	48 .006	194 .025	13 .002	
72000 *	9 .001	1 .000	36 .005	80 .010	
	* L	25	200	1000	5000
I = 17 J = 5	*				
K	*****	*****	*****	*****	*****
1000 *	95 .011	10 .C01	0 .000	0 .000	
6000 *	125 .016	325 .042	64 .00F	0 .000	
30000 *	0 .000	49 .006	366 .047	56 .007	
72000 *	0 .000	1 .000	30 .004	165 .021	
	* L	25	200	1000	5000
I = 17 J = 11	*				
K	*****	*****	*****	*****	*****
1000 *	0 .000	0 .C00	0 .000	0 .000	
6000 *	11 .001	139 .C18	34 .004	0 .000	
30000 *	0 .000	15 .002	474 .061	109 .014	
72000 *	0 .000	0 .000	67 .009	501 .064	

23

I - CORE
J - FILES
K - CPU
L - IO

	* L	25				200				1000				5000			
		*****				*****				*****				*****			
I= 29	*	*****				*****				*****				*****			
J= 3	K	*****				*****				*****				*****			
	*	*****				*****				*****				*****			
	1000 *	84	.011	10	.001	0	.000	0	.000	0	.000	0	.000	0	.000	0	.000
	6000 *	11	.001	28	.004	2	.000	0	.000	0	.000	0	.000	0	.000	0	.000
	30000 *	38	.005	161	.021	53	.007	0	.000	0	.000	0	.000	0	.000	0	.000
	72000 *	12	.002	3	.000	56	.007	17	.002	17	.002	17	.002	17	.002	17	.002

	* L	25				200				1000				5000			
		*****				*****				*****				*****			
I= 29	*	*****				*****				*****				*****			
J= 5	K	*****				*****				*****				*****			
	*	*****				*****				*****				*****			
	1000 *	7	.001	0	.000	0	.000	0	.000	0	.000	0	.000	0	.000	0	.000
	6000 *	10	.001	61	.008	6	.001	0	.000	0	.000	0	.000	0	.000	0	.000
	30000 *	1	.000	14	.002	56	.007	0	.000	0	.000	0	.000	0	.000	0	.000
	72000 *	1	.000	11	.001	47	.006	62	.006	62	.006	62	.006	62	.006	62	.006

	* L	25				200				1000				5000			
		*****				*****				*****				*****			
I= 29	*	*****				*****				*****				*****			
J= 11	K	*****				*****				*****				*****			
	*	*****				*****				*****				*****			
	1000 *	0	.000	0	.000	0	.000	0	.000	0	.000	0	.000	0	.000	0	.000
	6000 *	3	.000	17	.002	0	.000	0	.000	0	.000	0	.000	0	.000	0	.000
	30000 *	0	.000	7	.001	70	.009	2	.000	2	.000	2	.000	2	.000	2	.000
	72000 *	0	.000	0	.000	33	.004	125	.016	125	.016	125	.016	125	.016	125	.016

Table IV -- Joint Probability Distribution of the Real Workload (Continued)

to disk. Seventy-five calibration runs were conducted to relate these two parameters to the direct time and the block count. In all these calibration runs the LOGGER summary data was used to determine the block count and the direct time. Various combinations of the four parameters, viz., core size, number of files, and the parameters for direct time and the block count were used in these calibration experiments. Table V presents these experimental results.

These experimental results are used in the inversion of the test workload characteristics into synthetic program parameters. During the calibration experiments, the synthetic program parameters are varied over the required range and the resulting values of the direct time and the number of block counts are obtained from the LOGGER summary data. Inversion consists of reversing the above procedure. In other words, given a set of workload characteristics, (direct time and the number of block counts) the required values of the synthetic program parameters are determined from the calibration results. In this study because of the number of variables involved, no attempt was made to derive general expressions relating the synthetic program parameters with the workload characteristics. Instead the calibration runs were conducted with the values of the parameters in the immediate neighborhood of the desired workload characteristics. The experimental results were used to determine the parameter settings for the 87 jobs in the test workload. The synthetic workload is the collection of these 87 jobs.

CLOSURE

The performance of a computer system can be described in terms of the interaction between the workload and the hardware-software configuration. A method of studying this interaction is to conduct experiments with the existing hardware-software configuration using a stabilized, reproducible workload that is representative of the

Table V
Experimental Results Used in Calibrating the Synthetic Program

Job No.	Synthetic Program Parameters				Calibration Results			
	Compile Time		Run Time		Core (K-bytes) (i)	Files (j)	CPU (sec) (k)	Block Count (l)
	Core (i)	Files (j)	Times thru CPU loop (k)	Total Block Count (l)				
1	4	3	100	25	4.5	3	1	25
2	4	3	100	200	4.5	3	3	200
3	4	3	300	1000	4.5	3	11	1000
4	4	6	1000	200	4.5	6	11	200
5	4	6	800	5000	4.5	6	49	5000
6	4	11	3000	1000	6.5	11	37	1000
7	17	3	100	25	17.5	3	2	25
8	17	3	300	1000	17.5	3	11	1000
9	17	3	10000	5000	17.5	3	139	5000
10	17	6	600	25	17.5	6	8	25
11	17	6	500	1000	17.5	6	13	1000
12	17	6	4000	5000	17.5	6	82	5000
13	17	11	800	200	18.0	11	13	200
14	17	11	1000	5000	18.0	11	41	5000
15	17	11	4000	5000	18.0	11	89	5000
16	29	3	100	25	29.5	3	6	25
17	29	3	2500	200	29.5	3	28	200
18	29	3	5000	25	29.5	3	46	25
19	29	6	10	200	29.5	6	2	200
20	29	6	800	200	29.5	6	12	200
21	29	6	4000	5000	29.5	6	86	5000
22	29	11	4000	5000	30.0	11	84	5000
23	29	11	15000	1000	30.0	11	146	1000
24	4	3	100	25	4.5	3	2	25
25	4	3	50	200	4.5	3	4	200
26	4	3	300	200	4.5	3	5	200
27	4	3	50	1000	4.5	3	9	1000
28	4	3	2500	1000	4.5	3	37	1000
29	4	3	3200	5000	4.5	3	77	5000
30	4	6	300	200	4.5	6	6	200
31	4	6	50	1000	4.5	6	10	1000
32	4	6	2500	1000	4.5	6	36	1000
33	4	6	100	5000	4.5	6	45	5000
34	4	11	2500	1000	6.5	11	38	1000
35	4	11	3000	5000	6.5	11	72	5000
36	17	3	70	25	17.5	3	1	25
37	17	3	100	200	17.5	3	3	200

Table V
Experimental Results Used in Calibrating the Synthetic Program
(Continued)

Job No.	Synthetic Program Parameters				Calibration Results			
	Compile Time		Run Time		Core (K-bytes) (i)	Files (j)	CPU (sec) (k)	Block Count (l)
	Core (i)	Files (j)	Times thru CPU loop (k)	Total Block Count (l)				
38	17	3	10	1000	17.5	3	10	1000
39	17	3	2000	1000	17.5	3	32	1000
40	17	3	4000	5000	17.5	3	76	5000
41	17	6	10	25	17.5	6	2	25
42	17	6	400	25	17.5	6	6	25
43	17	6	100	200	17.5	6	4	200
44	17	6	10	1000	17.5	6	9	1000
45	17	6	2000	1000	17.5	6	30	1000
46	17	6	3000	5000	17.5	6	72	5000
47	17	11	100	200	18.0	11	5	200
48	17	11	2300	1000	18.0	11	33	1000
49	17	11	500	5000	18.0	11	47	5000
50	17	11	6500	1000	18.0	11	76	1000
51	17	11	3000	5000	18.0	11	75	5000
52	29	3	10	25	29.5	3	2	25
53	29	3	2600	200	29.5	3	28	200
54	29	6	400	200	29.5	6	7	200
55	29	6	3000	1000	29.5	6	48	1000
56	29	6	3000	5000	29.5	6	74	5000
57	29	11	3000	1000	30.0	11	38	1000
58	29	11	3000	5000	30.0	11	72	5000
59	4	3	0	200	4.5	3	4	200
60	4	3	0	1000	4.5	3	9	1000
61	4	3	2000	1000	4.5	3	29	1000
62	4	3	3000	5000	4.5	3	73	5000
63	4	6	0	1000	4.5	6	9	1000
64	4	6	2000	1000	4.5	6	31	1000
65	4	6	0	5000	4.5	6	43	5000
66	4	11	2000	1000	6.5	11	33	1000
67	17	3	400	200	17.5	3	6	200
68	17	3	0	1000	17.5	3	9	1000
69	17	6	350	200	17.5	6	7	200
70	17	6	0	1000	17.5	6	11	1000
71	17	11	300	200	18.0	11	7	200
72	17	11	2000	1000	18.0	11	34	1000
73	17	11	0	5000	18.0	11	41	5000
74	29	6	2000	1000	29.5	6	30	1000
75	29	11	2000	1000	30.0	11	29	1000

real workload. That there is a need for constructing a representative test workload need not be overemphasized. In this study the representative workload is constructed by first isolating the most frequently occurring jobs.

The accounting package seldom captures data about the human factors involved, e.g., time for mounting and dismounting tapes. Special measurements have to be made to determine the human factors. Typical measurements are the number of frequently used tapes, and the method of tape assignments.

The following procedure may be adopted for evaluating the performance of a computer system in which tape files dominate. The workload can be divided into two distinct classes - jobs with disk I/O only and jobs with tape and/or disk I/O. The former class can be represented using methods described in this section. The latter class can be represented with the use of synthetic programs that create tape files. The combined representative synthetic workload can then be used to study the effects of system parameters; for example, changes in the configuration, addition to the existing configuration, blocking factor etc. In these experiments care should be taken to insure that the human factors are adequately simulated by introducing known amounts of delay in mounting and dismounting tapes.

SECTION III

EXPERIMENTS WITH THE TEST WORKLOAD

PURPOSE

A computer system can be considered as a hardware-software-configuration (HSC) interacting with its workload. Evaluation of computer system performance requires an understanding of this interaction in order to assess and improve the performance and to predict the effects of changes in either the workload or the HSC.

In this study, experiments were conducted with the B 3500 using a representative model of the real workload. The controlled experiments were conducted in a dedicated environment. A hardware monitor (DYNAPROBE) was used to measure the utilization of the processor, the channel and the physical devices.

EXPERIMENTAL RESULTS

The test workload, representative of the batch jobs processed during a month on the B 3500 (ESD, Hanscom Air Force Base) consisted of 87 jobs. Their CPU time, core used, number of files and block count to disk are presented in Table III. LOGGER data was used to obtain the following characteristics for each job.

- a) Start Time
- b) Stop Time
- c) Direct time used
- d) Core requested
- e) Number of files
- f) Block count per file

A hardware monitor, DYNAPROBE 7900, was used to measure the following quantities.

- a) CPU busy in the normal state
- b) CPU idle
- c) Disk channel 2 busy (primary)
- d) Disk channel 10 busy (alternate)
- e) Disk channels 2 and 10 busy

Figure 4 is a schematic of the hardware configuration used in these experiments. The electrical signals corresponding to the above quantities were recorded on tape which was later analyzed to determine the utilization values.

The architecture of the B 3500 was initially studied to determine the nature of the experiments to be performed. In this system there are six disk units connected to the processor by a pair of channels (channel 2 and 10) to achieve simultaneity. The disk units are head-per-track units. As there is only one primary channel (channel 2) for all the disk units, overlap between data transfers can be accomplished by forcing the transfer to take place through the secondary channel (channel 10). Because of this channel - disk unit relationship it is not possible to initiate many data transfers, analogous to parallel seeks characteristic of moving arm disk units.

The degree of multiprogramming, the number of jobs co-resident in main memory, is determined mainly by the core size of the available jobs in the mix. In the experiments reported here, all the jobs were initially spooled to the disk. The first set of jobs were initiated. The number of jobs in this set was determined by the individual core sizes. At the termination of a job, the operating system scans the list of available jobs and initiates a job which can fit into the vacant core space. If such a job cannot be found the operating system waits for a second job to terminate and the search for




Figure 4. B-3500 HARDWARE CONFIGURATION

the next job is repeated with the difference that the vacant space now available has increased. For this reason the degree of multiprogramming varies during a session. The multiprogramming (number of jobs in the mix) resulting from one of the experimental sessions is shown in Figure 5.

Eight runs were conducted using the test workload. The effect of external sequencing (i.e., the order in which the jobs are processed by the system) on the overall performance of the computer system was studied. The first experiment was run single-thread to verify the characteristics of the test workload. The second experiment was run with multiprogramming in which the jobs were initiated from the operator's console after every job termination. This was found to be inefficient as the system was idling waiting for jobs to be initiated. In the third experiment the job initiating was accomplished automatically by the operating system with no operator intervention. It was found that changes in the external sequencing have no appreciable effect on the overall performance. This is not very surprising in view of the fact that the workload is CPU-bound. In single processor, multiprogrammed systems, CPU-bound workloads lead to a situation in which all the jobs have to wait to use the processor, since the overlap between the processor and I/O usage is very small.

The test workload used only disk files and the I/O activities were, therefore, much faster than the equivalent I/O activities to tape. In effect, the I/O load on the system was reduced and the test workload was made CPU-bound. This is confirmed by the fact that overall CPU utilization of the month's workload studied was approximately 25% and the CPU utilization of the test workload was approximately 65%. It was necessary to adjust the I/O load on the system so that the CPU utilization became reasonable.




Figure 5. VARIATION OF THE DEGREE OF MULTIPROGRAMMING DURING EXPERIMENT 4

The synthetic workload used allows for adjustment of the I/O load by means of the I/O run time program parameter. The assumption is made that increasing the I/O load will bring the CPU utilization down to a value approaching that observed for the real workload. This implies that the neglected tape I/O activities can be replaced by a suitable number of disk I/O activities. The I/O control parameter for all the jobs in the test workload was increased by a factor of 5. The result was that the elapsed time of the experiment became too large. Only 36 of the total of 87 jobs were completed in slightly more than 2 hours. The number of block counts was then increased by a factor of 2 and the resulting processor utilization was found to be 45.8%. In both these experiments it was found that replacing the tape I/O with disk I/O led to large values of channel utilization. In trying to decrease the CPU utilization to realistic values we were in effect increasing the channel utilization beyond the realistic values. In real life situations the tape drives and disk drives are connected through separate channels and it is not possible to study the behaviour of such a system by using a test workload that places the total I/O load on any one of the channels. The overall summary of experimental results is shown in Table VI.

HARDWARE MONITORING

A hardware monitor, DYNAPROBE 7900*, is used to measure the resource utilization of the major resources during the controlled experiments. The hardware monitor is a high impedance meter that measures the electrical signals that correspond to the busy or the idle state of a given resource. The hardware monitor, because of its high impedance, does not perturb the system being measured. This is a decided advantage over software monitors which influence the performance of the system being measured. The hardware monitor

*Comress, Incorporated.

Table VI
Overall Summary of the Experimental Results

Experiment Number	Elapsed Time (sec.)	Processor Time (sec.)	Total Block Count	CPU Utilization	Remarks
1	6896	1794	110,775	26.0%	87 jobs, single thread, structured schedule.
2	3344	1994	110,750	59.6%	86 jobs, multi-programming by executing jobs from operator's console; structured schedule.
3	3025	2034	110,750	67.2%	86 jobs, multi-programming with MCP scheduling and executing the jobs; structured schedule.
4	3682	2071	110,775	56.2%	87 jobs, multi-programming (same as Exp. 3), unstructured schedule.
5	7726	2669	216,733	34.5%	36 jobs, 8 partial jobs, multi-programming (same as Exp. 3), unstructured schedule, increased I/O by factor of 5.
6	5825	2666	202,100	45.8%	74 jobs, incomplete run, multi-programming (same as Exp. 3), unstructured schedule, doubled I/O activity.
7	6396	2937	221,550	45.9%	87 jobs, same as Exp. 6
8	5127	2936	221,550	57.3%	87 jobs, multi-programming (same as Exp. 3), structured schedule, doubled I/O activity.

captures the state of the various system resources continuously and its output is written out to tape. The hardware monitor is capable of measuring the state of a large number of system resources simultaneously, by means of hardware probes that are connected to the appropriate pins in the computer. These probe signals can be combined (logically) to determine the degree of overlap. The hardware monitor output is collected and later analyzed by means of special software routines to determine the utilization values.

The following were measured using the hardware monitor.

- 1) Processor busy in the normal state
- 2) Processor idle
- 3) Channel 2 busy
- 4) Channel 10 busy
- 5) Channel 2 and channel 10 busy

At the end of the experiment the output tape is analyzed by a special software package. The output of this program appears in Table VII. The program calculates two sets of utilization for the resources indicated in the table. The first value is the value for the interval. The second value is the cumulative value for the duration of the test period. The cumulative values of the utilization are used for the purpose of comparison with the corresponding values calculated from the accounting data.

COUNTER	DESCRIPTION	(BASE = T01)	TOTALS FOR LAST	51.0 SECONDS	ACCU. TOTALS FOR	521.0 SECONDS
C01	NORMAL STATE		3.3375494 SEC	6.56 PCT	319.4314924 SEC	61.31 PCT
C02	MCP IOLE LOOP		46.2411553 SEC	92.46 PCT	1774.3329957 SEC	36.11 PCT
C03	CHANNEL 2 BUSY PRIMARY		1.022693 SEC	.21 PCT	4317.2165858 SEC	77.11 PCT
C04	CHANNEL 10 BUSY ALTERNATE		1.030300 SEC	.21 PCT	3573.2103433 SEC	73.60 PCT
C05	CHANNELS 2 AND 10 BUSY OVERLAP		1.030300 SEC	.21 PCT	2743.5481359 SEC	53.62 PCT
C06	E01 1ST 50K ADDRESSES BUSY		1.022052 SEC	.20 PCT	422.3815905 SEC	17.72 PCT
C07	E02 50K ADDRESSES BUSY		1.030300 SEC	.21 PCT	101.3834765 SEC	2.53 PCT
C08	E03 3RD 50K ADDRESSES BUSY		1.030300 SEC	.21 PCT	144.1744987 SEC	2.86 PCT
C09	E04 4TH 50K ADDRESSES BUSY		1.030300 SEC	.21 PCT	2073.3598252 SEC	57.19 PCT
C10	E01 M00 1 ACTIVE TIME		1.030300 SEC	.21 PCT	147.1171334 SEC	3.59 PCT
C11	E01 M00 2 ACTIVE TIME		1.030300 SEC	.21 PCT	273.272371 SEC	5.34 PCT
C12	E01 M00 3 ACTIVE TIME		1.030300 SEC	.21 PCT	7.7163456 SEC	1.28 PCT
C13	E01 M00 4 ACTIVE TIME		1.030300 SEC	.21 PCT	7.7163456 SEC	1.28 PCT
C14	E01 M00 5 ACTIVE TIME		1.030300 SEC	.21 PCT	7.7163456 SEC	1.28 PCT
C15	E02 WITH CHANNEL 2 BUSY		1.030300 SEC	.21 PCT	2475.5643122 SEC	57.11 PCT
C16	E01 WITH CHANNEL 2 BUSY		1.030300 SEC	.21 PCT	259.3493743 SEC	44.91 PCT
T01	TOTAL ELAPSED TIME		5.1197180 SEC	14.73 PCT	1453.8665753 SEC	28.19 PCT
T02	PAUSE STATE		1.030300 SEC	.21 PCT	3213.111300 SEC	100.00 PCT
T03	STOP STATE		1.030300 SEC	.21 PCT	1.030300 SEC	0.00 PCT
C01	CPU ACTIVE XNDXAL STATE & MCP		2.7389447 SEC	7.52 PCT	3431.5571743 SEC	65.87 PCT
C02	CONTROL STATE XNDXAL STATE & I/O CONTROL		1.030300 SEC	.21 PCT	257.5355135 SEC	4.56 PCT
C03	CHANNEL 2 OVERLAP WITH CPU ACTIVE		2.7389447 SEC	.21 PCT	2543.1163052 SEC	51.44 PCT
C04	CHANNEL 2 OVERLAP WITH CPU IOLE		1.030300 SEC	.21 PCT	337.1193925 SEC	23.66 PCT
C05	CHANNEL 10 OVERLAP WITH CPU ACTIVE		1.030300 SEC	.21 PCT	231.9738247 SEC	48.60 PCT
C06	CHANNEL 10 OVERLAP WITH CPU IOLE		1.030300 SEC	.21 PCT	1145.2363752 SEC	22.38 PCT
C07	CHANNEL 2 OR 10 BUSY		1.030300 SEC	.21 PCT	491.1396839 SEC	98.09 PCT
C08	CHANNEL 2 OR 10 OVERLAP WITH CPU ACTIVE		2.7389447 SEC	.21 PCT	3271.7247935 SEC	62.86 PCT
C09	CHANNEL 2 OR 10 OVERLAP WITH CPU IOLE		1.030300 SEC	.21 PCT	1512.1711155 SEC	31.29 PCT
C10	CONTROL STATE OVERLAP WITH E01 1ST FASE		1.030300 SEC	.21 PCT	534.1827654 SEC	12.17 PCT
C11	E03 OVERLAP WITH SYSTEM ACTIVE		1.030300 SEC	.21 PCT	1743.7112935 SEC	33.55 PCT
C12	E03 OVERLAP WITH SYSTEM IOLE		1.030300 SEC	.21 PCT	31.12782955 SEC	15.36 PCT
C13	E01 OVERLAP WITH SYSTEM ACTIVE		1.030300 SEC	.21 PCT	332.3455691 SEC	17.49 PCT
C14	E01 OVERLAP WITH SYSTEM IOLE		1.030300 SEC	.21 PCT	73.4241323 SEC	1.35 PCT
C15	SYSTEM ACTIVE XNDXAL CPU & 2 OR 10 BUSY		2.6345159 SEC	7.70 PCT	3051.7351135 SEC	97.15 PCT
C16	SYSTEM IOLE		4.184864 SEC	92.46 PCT	245.2513822 SEC	2.85 PCT
SYSTEM UTILIZATION XNDXAL WITH CHANNEL 2 OR 10						
FROM 23.25.15.0 TO 23.26.00.0						
DESCRIPTION / PERCENT OF T01						
XXXXXXXXXX						
SYSTEM ACTIVE						
SYSTEM IOLE						
NORMAL STATE USER PGMS ONLY XXXXXXXX						
CPU OVERHEAD						
2/10 AND CPU ACTIVE						
2/10 AND CPU IOLE						

Table VII --- DYNAPROBE Hardware Monitor Reduction Output

SECTION IV

SUMMARY

Computer performance studies can be classified into two groups - assess and improve the performance; and predict the effects of changes in either the workload or the hardware-software configuration. Efforts in this area are greatly helped by a proper understanding of the interaction between the workload and the hardware-software-configuration. This report presents an approach that may be useful in understanding this interaction.

Experiments are performed with a Burroughs B 3500 computer system (ESD, Hanscom Air Force Base) using a synthetic test workload, in a dedicated environment. The method of constructing the test workload has been described. The method consists of matching the joint probability density of the real workload with that of the test workload. Direct time, I/O to disk, core size and the number of files are the four characteristics selected for representation. The magnitudes of these characteristics are derived from the LOGGER summary data.

In these controlled experiments the resource utilization is monitored using a hardware monitor. The values of the utilization calculated from the accounting data compare favorably with those measured by the monitor. This is a very significant finding of this study. This study has shown that the accounting data collected by the LOGGER is sufficient to calculate the processor and the channel utilization for the B 3500 computer system. LOGGER provides a ready and economical source of data for calculating the resource utilization that is very useful in analyzing the performance of the B 3500 system. A system manager can conveniently assess the system performance by periodic processing of the accounting data to determine the resource utilization.

The test workload used in these experiments does not include tape I/O activities. In installations where the tape I/O activities dominate, the synthetic workload experiments should be preceded by an analysis of the human factors involved. Some examples of these human factors are the layout of the tape library, the number of tapes in the library and the time taken for mounting and dismounting tapes. These human factors have a significant influence on the computer performance.

The tape I/O can be accounted for by creating tape files in the synthetic program. The workload can be divided into two distinct classes - jobs with disk I/O only and jobs with tape and/or disk I/O. The combined synthetic workload can then be used to study the effects of system parameters; for example, changes in the configuration, addition to the existing configuration, blocking factor etc.

Representative, synthetic workloads are not ends in themselves but just means to an end. They can be used as stabilized, reproducible workloads in conducting experiments with an existing hardware-software-configuration to evaluate the performance of the computer system. The synthetic workloads have the additional advantage that they are flexible. By varying the synthetic program parameters the characteristics of the test workload can be changed.

APPENDIX

CALCULATION OF THE RESOURCE UTILIZATION

The utilization of the processor and the I/O channels was monitored with a hardware monitor. The accounting package, LOGGER, collected the direct time for each job and the total number of block counts transferred by each job. In this section methods for calculating the resource utilization from the accounting data are described. The calculated values of the resource utilization are then compared with those measured by the hardware monitor.

PROCESSOR

The direct time (d_i) of every job in the synthetic workload is determined from the LOGGER data. Let the elapsed time for the synthetic workload be T and N be the total number of jobs in the workload. ($N = 87$ in our study.) The processor utilization $\bar{O}_1$ in the normal state is given

$$\bar{O}_1 = \frac{1}{T} \sum_{i=1}^N d_i \quad (A1)$$

The processor utilizations for the seven runs were calculated using the equation (A1). Table VIII presents these values and the corresponding values measured by the hardware monitor are also presented in the table for the purposes of comparison. The two sets of values agree fairly well. The LOGGER itself consumes some CPU time but it does not measure itself. The hardware monitor on the other hand, measures the LOGGER activity. The time for interrupt processing is usually charged to the interrupted job if no job switching

Table VIII
Comparison of the Processor Utilization Values

Experiment Number	Processor Utilization (%)	
	Accounting Data	Hardware Monitor
1	26.0	—
2	59.6	62.0
3	67.2	64.1
4	56.2	60.2
5	34.5	34.5
6	45.8	48.2
7	45.9	50.9
8	57.3	61.3

takes place. The LOGGER collects the data whenever the processor is in the normal state and does not collect the data when the system is in the master state (i.e., supervisory state). But the hardware monitor does collect the data for both the states. These considerations should be borne in mind when comparing the results in Table VIII.

I/O CHANNEL

The LOGGER records the total number of block counts to each channel. In the channel-device architecture, there is a primary channel (channel 2) and a secondary channel (channel 10) to achieve simultaneity. This results in the transfer of the block by channel 2. Whenever it is busy channel 10 takes over the transfer. The LOGGER does not distinguish between the number of block counts transferred to the primary and the secondary channel, but measures the total block count for the primary and the secondary combination. The hardware monitor, on the other hand, distinguishes between the primary and the secondary channel and measures the utilization of the two channels separately.

The devices used in the B3500 system are head-per-track disks and the data transfer takes place in three phases. During the first phase the track address is analyzed and the corresponding read/write head is connected to the channel. During the second phase the read/write head waits for the beginning of the track to rotate. The second phase is referred to as the latency. During the third phase the data transfer takes place. It should be pointed out that the channel and the device are busy during all the three phases. The seek is absent in the transfer process since no arm movement is involved.

The time per block count, t , can be expressed as

$$t = (\text{ave. Latency}) + (\text{Transfer Time}) + (\text{Overhead}) \quad (\text{A2})$$

where the average latency is the time for half a revolution of the disk, the transfer time is a function of the block size and the rate of transfer and the overhead includes the time for switching from track to track. The hardware characteristics of the disk-channel unit are given below.

<u>Type</u>	B9372-9
Rotation Speed	1300 RPM
Max. Latency	46 milliseconds
Ave. Latency	23 milliseconds
Rate of Transfer	377 kilobytes/second
Overhead	1.8 milliseconds

The block size for all the files in the test workload was 50 bytes. Based on these values the average time per block count is given by

$$t = 23.0 + \frac{50}{377} + 1.8 \approx 23.0 + 0.13 + 1.8 \approx 25 \text{ milliseconds} \quad (\text{A3})$$

The channel utilization can be calculated from the following expression

$$\mathcal{U}_2 = \frac{(t)(B)}{T} \quad (\text{A4})$$

where

$\mathcal{U}_2$ - channel utilization

t - time per block count = 0.025 seconds

B - total block count

T - elapsed time for the test workload

REFERENCES

1. ADPE Performance Management System, AFM 171-400, Volume III, 1 December 1972.
2. Burroughs B 2500 and B 3500 Information Processing Systems Assembler Reference Manual, PCN 1034949, November 20, 1970.
3. J. Esposito, "Statistical Analysis to Determine Digital Computer Workload Characteristics," ESD-TR-74-175, June 1974.
4. K. Sreenivasan and A. J. Kleinman, "Construction and Applications of Representative Synthetic Workloads," ESD-TR-73-212, August 1973.