MICROCOPY RESOLUTION TEST CHART

NATIONAL BUREAU OF STANDARDS-1963-A
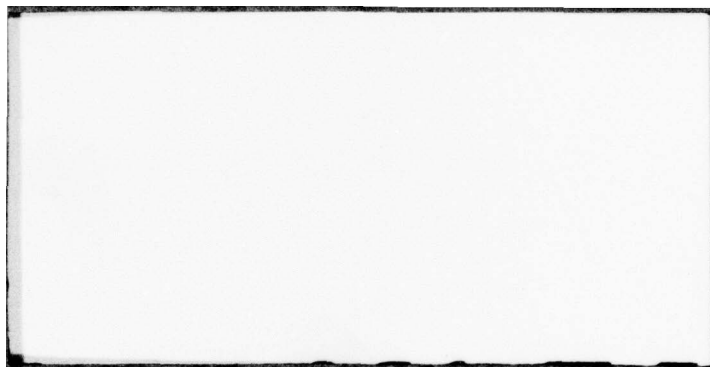
AFIT/GCS/EE/78-20

LEVEL H

①

⑥ COMPUTER PERFORMANCE EVALUATION
OF INDIVIDUAL WORKLOADS ON
THE ASD CDC CYBER SYSTEM.

⑨ Master's THESIS,

⑭ AFIT/GCS/EE/78-20

⑩ Jimmy W. Thompson
Capt   USAF

⑪ Dec 78   ⑫ 96p.

Approved for public release; distribution unlimited

012 225   LB

79 01 30 108

COMPUTER PERFORMANCE EVALUATION

OF INDIVIDUAL WORKLOADS ON THE ASD

CDC CYBER SYSTEM

THESIS

Presented to the Faculty of the School of Engineering

of the Air Force Institute of Technology

Air University

in Partial Fulfillment of the

Requirements for the Degree of

Master of Science

by

Jimmy W. Thompson, B.S.

Capt   USAF

Graduate Computer Systems

December 1978

Preface

This thesis is an investigation into the applications workload imposed on a Control Data Corporation Cyber 74 computer system, operated by the Air Force Aeronautical Systems Division at Wright-Patterson AFB, Ohio. Once the workload was defined, the resulting service delivered to the organizational users was calculated and analyzed. The methodology and results yielded some useful conclusions concerning the Cyber 74 operation, and a background for further computer performance evaluation and measurement studies.

There are many people who aided in the completion of this thesis. Dr. Ed Reeves and Dr. Gary Lamont provided able assistance as committee members. Special thanks go to my advisor, Dr. Tom Hartrum. Without his personal interest on my behalf, this thesis would have gone uncompleted.

For providing emotional support above and beyond the call of duty, I would like to thank my wife, Rosemary.

<div align="right">Jimmy W. Thompson</div>

Contents

List of Figures

List of Tables

Abstract

This paper describes computer performance evaluation and measurement techniques applied to a Control Data Corporation Cyber 74 system, located at Wright-Patterson AFB, Ohio. The performance of a computer system is influenced by its hardware configuration, system software, man-machine interaction, and application program workload. The focus of this investigation is to determine how application workload affects performance. Application workload is measured by the magnitude of job demands placed on the system hardware resources; for example, CPU and I-O time, and central memory used. System accounting data provides resource usage values as well as related delay times for each job. Workload is then characterized by statistics for each job class and by total job statistics. The grouping of jobs into classes is done using clustering, a multivariate data analysis technique. Batch turnaround time is then analyzed for its dependence on workload. The methodology used and the results obtained provide performance conclusions and a background for further research.

COMPUTER PERFORMANCE EVALUATION OF INDIVIDUAL

USER WORKLOADS ON THE CDC CYBER SYSTEM

I. Introduction

Background

Computer performance evaluation and measurement (CPEM) can be
undertaken for a variety of reasons. Examples include the selection
of a new computer, the design of new hardware, and the improvement of
existing systems (Ref 24: 3-4). This investigation focuses on per-
formance evaluation of an existing system. A fundamental problem of
computer performance analysis is the definition of performance and a
determination of criteria that control performance. For this investi-
gation, performance is defined as the effectiveness with which a computer
uses its resources to accomplish system objectives (Ref 23: 8). These
objectives can be discussed from two viewpoints: first, user considera-
tions such as turnaround time; and second, system efficiency such as
throughput and CPU utilization. Table I contains a definition of these
performance measures.

Having selected performance measures, the crucial problem is to
determine how these measures depend on the system workload and the system
design. An understanding of such a relationship is essential if perform-
ance optimization efforts are to be constructive (Ref 23: 15-19). Cer-
tainly, hardware features such as disk access techniques and tape drive
speed impact performance. Also affecting performance is the computer
operating system, part of which creates an overhead workload. The final

1

criteria effecting performance is the demand on system resources caused by the execution of the application programs.

This investigation was based on the performance of a system as affected by its workload. For this effort, workload is defined as the amount of hardware resources consumed by the execution of application programs. User resource consumption is quantified by the use of job parameters as shown in Table II. Once the performance measures have been selected and the performance criteria specified, methods for quantifying performance for a given workload are developed.

The most accurate performance values are obtained when the system is measured under its actual workload. Because it is expensive and usually unreasonable to dedicate computer resources to continual performance measurement, techniques such as simulation or analytic performance models are often used (Ref 23: 19). Driving system performance models with the real workload is costly, time consuming, and inconvenient; therefore a need exists for a representative synthetic workload or a workload model (Ref 22:2). Examples of workload models are probabilistic, instruction mixes, benchmarks, and traces (Ref 23: 54-58). This investigation constructs a workload model using a cluster analysis technique (Ref 1). Then, system performance for the cluster workload model is obtained by empirical measurement. The specifics of workload modeling and performance measurement are explained in Chapter III.

Problem Statement

The Aerospace Systems Division (ASD), a component of the Air Force Systems Command, operates two large Control Data Corporation (CDC) computer

TABLE I

Performance Measures (Ref 23:16-17)

Measure	Description
Throughput	Amount of work completed per unit of time with a given workload
Turnaround time	Elapsed time between submitting a job to a system and receiving the output
CPU	Percent of time the CPU is doing useful work

TABLE II

Workload Characterization (Ref 23:12-13)

Workload Parameters	Description
Job CPU time	Total CPU time used by a single job
Job I/O	Total I/O time used by a single job
Central Memory	Maximum central memory used by a single job
Lines Printed	Number of lines printed by a single job

3

systems at Wright-Patterson Air Force Base, Ohio. The systems, a CDC 6600 and a Cyber 74, provide scientific processing support for ASD, a variety of Air Force research laboratories, Air Force contractors, and the Air Force Institute of Technology (AFIT). The support provided to these organizations is a combination of CDC 6600 and Cyber 74 processing.

The purpose of this investigation is to perform an analysis of the Cyber system to determine performance for individual organizations. Specific objectives are to characterize the Cyber workload on an organizational basis and to relate this workload to performance. The accomplishment of these objectives is limited by several constraints.

Constraints

A large amount of accounting data, in the form of job resource consumption and time parameters, is collected by the Cyber operating system, as "Dayfile" accounting data (Ref 9: 1-3). These parameters, needed for both workload and performance measurements, are collected daily in an event oriented format. The ASD computer center, as sponsors of this investigation, impose the following constraints on their support:

- -No modification to the operating system or accounting routines (Dayfile) would be allowed.

- -Additional processing of the Dayfile is inadvisable and no assistance could be given for obtaining data from the Dayfile.

Because of this constraint, workload and performance data was only available through the use of the Computer Load and Resource Analysis (CLARA) program (Ref 9). CLARA is used to reduce Dayfile data to a more manageable size. Because the reduced CLARA tapes did not have data available on interactive response times, performance analysis was limited to batch

4

performance parameters. Further impacts on this research as a result of missing CLARA data, as well as additional details about CLARA will be discussed in Chapter III. Data constraints, other than CLARA, included limits on the types and amount of data used. Because the workload affecting performance is greatest during normal duty hours, the analysis data was limited to the 0800-1700 hours time frame. Also, because no logs were kept regarding the physical input or output of jobs at computer access windows, it was not possible to include man-machine interaction times in user performance measures, such as turnaround time. Therefore, all time delays are measured using computer clock units.

Determination of the investigation objectives and constraints, in conjunction with a survey of available CFE literature, made possible the formulation of an investigation approach. This approach with supporting research references will be presented in the following section.

Approach

The first task undertaken was the development of a processing method to select and further reduce candidate workload and performance parameters from the CLARA tapes. The accomplishment of this task required that an understanding of the Cyber hardware and software and CLARA be achieved. References 3, 4, 8, 9, 10, 11, and 12 were used to aid in this effort, which is reported on in Chapters II and III. The second task consisted of selecting the best job parameters that characterize workload. Available literature (Ref 1, 2, 9, 17, 19 and 23) provided some assistance in this regard. By using the correlation program available in the Statistical Package for the Social Sciences (SPSS) (Ref 21),

5

the interdependencies of candidate workload parameters were examined, which further aided the selection process. After the workload parameters were chosen, the subsequent construction of a representative workload model posed the next problem. Two different types of workload models used were descriptive statistics (Ref 6, 20), and a cluster analysis technique (Ref 1, 2, 15, 17, 18, 19). Chapter III is a complete description of the workload characterization techniques which produced Cyber models for the system and individual user workloads. The determination of how each workload model affected each performance indicator was obtained from empirical measurements. This yielded an empirical performance model associated with a specific workload. The workload and resulting performance for the system, and the individual organizations are reported on in Chapter IV. Chapter V contains the conclusions reached as a result of the investigation, and recommendations for CDC performance improvement actions. But first, the ASD Cyber environment is presented.

II. ASD Cyber 74 Environment

The purpose of this chapter is to explain the Cyber hardware and support software so that their effect on performance can be analyzed in later chapters. To provide a background for the workload and performance analysis, the Cyber system, as used to support two different modes of operation is discussed.

Cyber 74 Hardware

The Cyber 74 was designed to overlap as many processes as possible (Ref 11: 4). The system provides for multiprogramming by using small independent computers, pheripheral processing units (PPUs), to accomplish input-output while the main CPU executes the computation instructions. Figure 1 reflects the major components of the Cyber 74. There are 20 identical peripheral control processors associated with the Cyber at ASD. Each is independent of the other and has its own 4K of memory for programs constructed from a 64 instruction repertoire. All peripheral processing units have access to 131K 60 bit words of central storage, and to the peripheral channels (Ref 4: 3-1). With this capability, the PPUs act as system control processors as well as I/O processors. This permits the central processor to continue high speed computations while the PPUs do the slower I/O and supervisory operations. Not only do the PPUs perform all I/O required by central storage programs, they also contain the major portion of the operating system, called the monitor. PPU number 0 is designated as the unit to hold the monitor, which is discussed further in the next section. A second PPU is assigned to control the operators console.

7
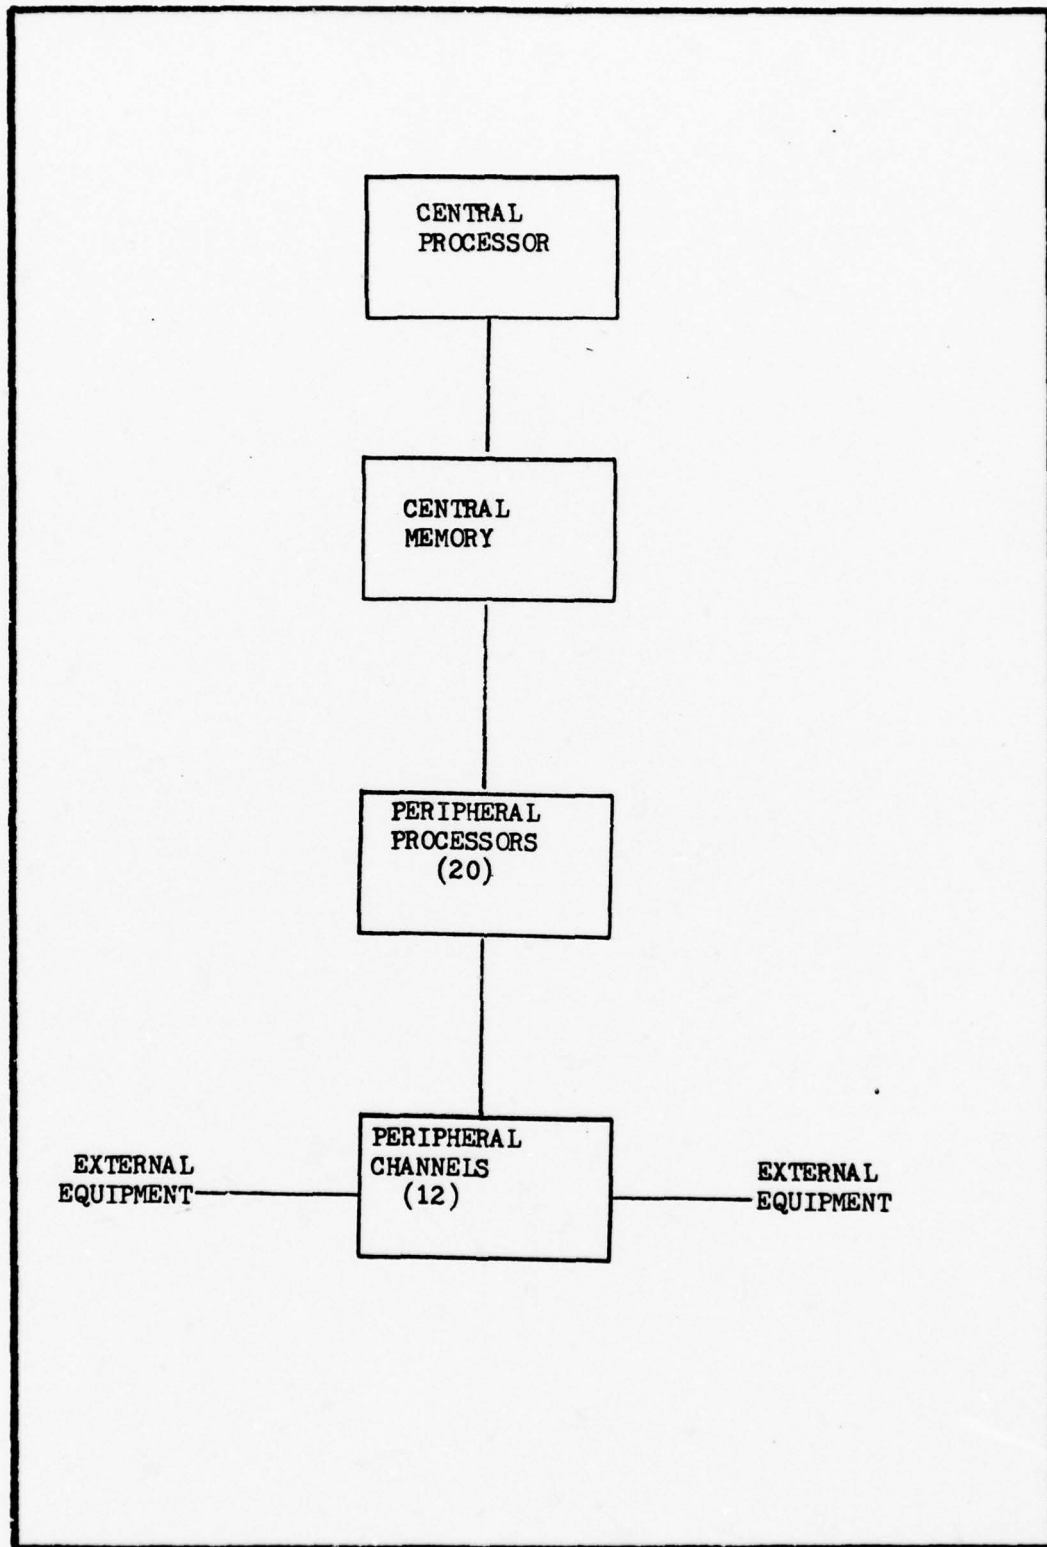
Figure 1. Cyber 74 Components

8

Cyber 74 Operation

In this section, the operating system and job scheduling algorithms are discussed to provide a background for later discussions of their effect upon performance. Critical to this discussion is the way the software impacts job delay time at each of the system queue points. First, the software is explained in general and then in more detail.

Overview. The Network Operating System/Batch Environment (NOS/BE) controls the operation of the Cyber 74 (Ref 4: 3-1). NOS/BE supports batch, interactive and graphics processing. An NOS/BE program JANUS, which is located in PPU 0 monitors the total operation of the system, including directing the CPU and other PPU actions (Ref 12: 2-1). To provide a multiprogramming capability, JANUS stores up to 15 jobs in central memory. These jobs reside in 15 variable length areas of memory called control points; the size of a control point is called its memory field length (Ref 13: 2-10). Additionally, JANUS is responsible for calling the scheduler program into execution to initiate job execution for new or swapped-out jobs (Ref 13: 12-13). The use of Scheduler will be discussed as it supports the flow of jobs through the system.

Job Flow. The job flow of the Cyber NOS/BE operating system is diagrammed in Figure 2. A job, which is a sequence of task and program steps, is read in and stored on disk. Those jobs initiated by control cards are stored in the input queue. Jobs spawned via intercom terminals are entered directly into the central memory (CM) queue. Also stored in the central memory queue are those jobs that have been swapped out (removed from a control point due to an expiration of their

9
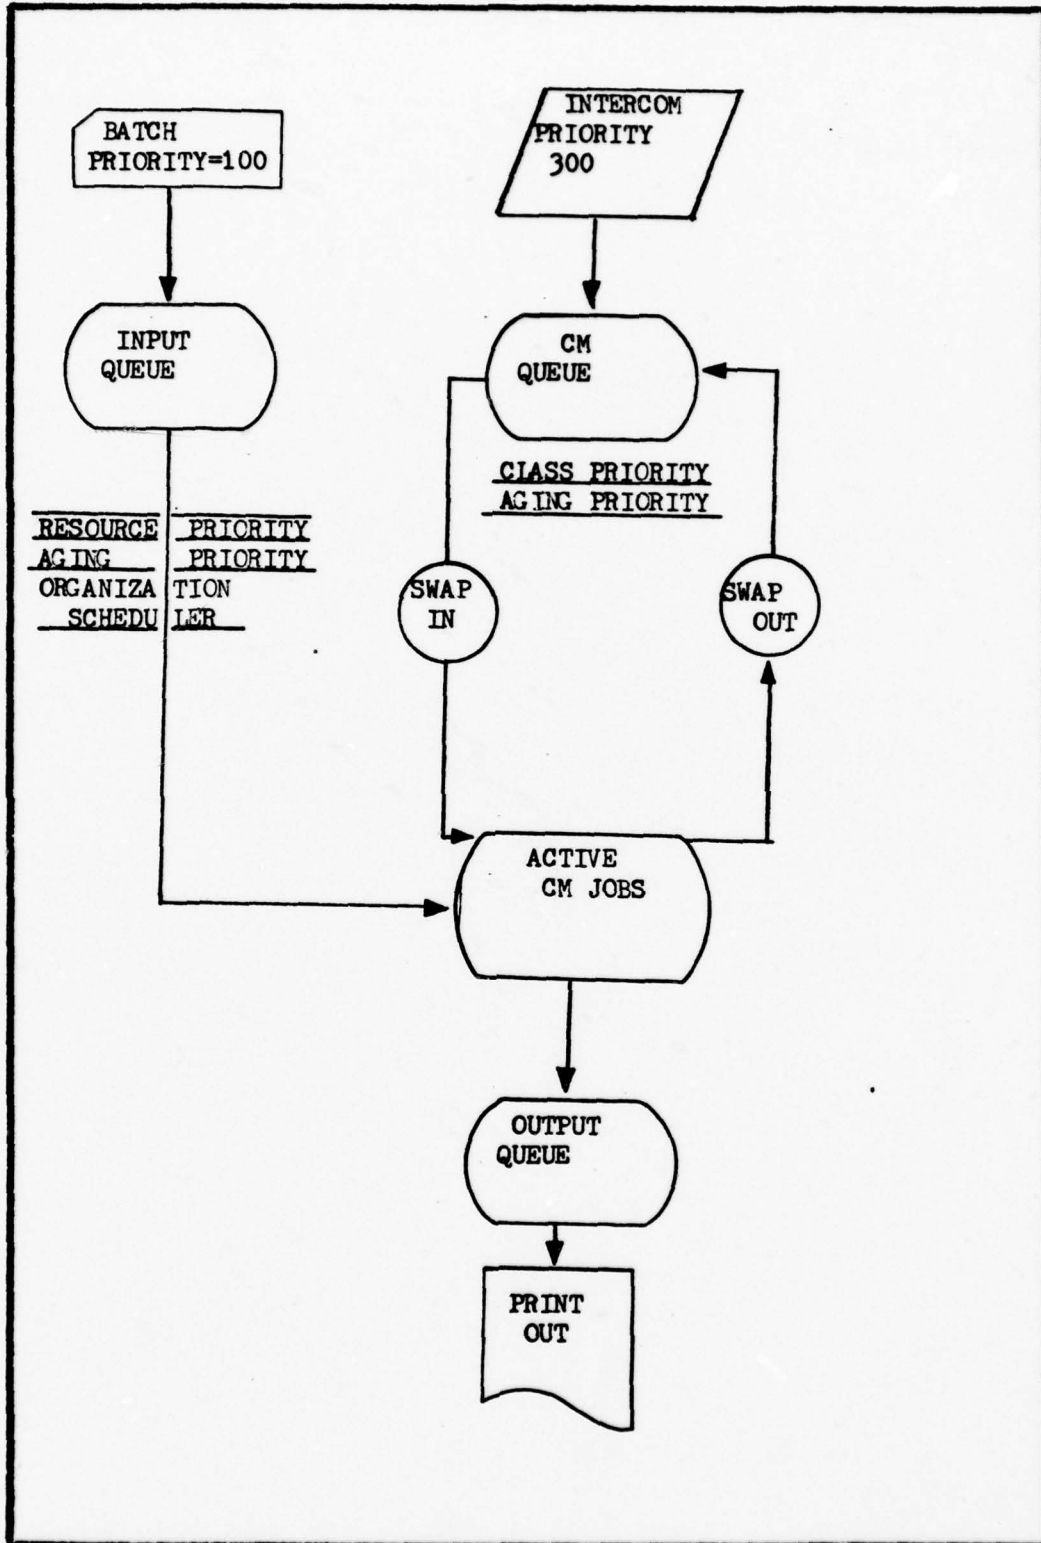
Figure 2. Cyber Job Flow(Ref 13: 1-2)

10

time slice). These jobs are all in competition for an active control point when it becomes available. Prior to determining how this contention can be resolved, some terms must be defined.

Job Classes are defined by CDC to provide a set of parameters for each class which can be used to compute queue priorities. The main job classes and their initial priorities are as follows (Ref 13: 8-9)

-Priority 1000 (Batch)

-Priority 3000 (Intercom)

-Priority 4000 (Multi-User), Currently only Editor.

-Priority 5000 (Express)

-Priority 6000 (Graphics)

Minimum Queue Priority (MINQP) is the priority given to a job in central memory which has used its allocated time slice (Ref 13: Chapter 1, 7).

Maximum Queue Priority (MAXQP) is the maximum priority level that a job in the CM queue may achieve while waiting for scheduling. All jobs in the CM queue are candidates for swap-in. Their priority in the CM queue continues to increase as the job ages, until it is swapped-in or MAXQP is reached (Ref 13: Chapter 1, 7).

Aging Rate (AR) is a factor used to weight the priority of a job according to the time it has spent in a wait queue (Ref 13: Chapter 1, 7).

Quantum Priority (QP) is the priority level assigned to a job once it has been swapped-in. The job will remain at this level until it has used its allotted time quantum (Ref 13: Chapter 1, 7).

11

<u>Base Quantum (BQ)</u> is a measure of time that a job will maintain a priority equal to Quantum Priority. The scheduler places more emphasis on CP time than PP time, so a ratio of 4 to 1 was chosen by CDC. That is, the quantum that any job has used is proportional to the amount of PP time (in sec) used.

The decision of whether the next job will come from the input queue or the central memory (CM) queue is dependent upon the circumstances that required the calling of Scheduler (Ref Figure 3). If a job has used its base quantum (BQ) and a higher priority job is in the CM queue, a swap-out and swap-in will be done. The priority of the jobs in the CM queue will never exceed the quantum priority (QP). That is MAXQP will always be less than QP to prevent unnecessary swapping.

Secondly, if a job has terminated, thereby giving up its control point memory, the Scheduler will check the input queue for a job to initiate. The highest priority job in the input queue will then be given a QP and brought into memory (Ref 8: 25-33). The priority of jobs in the input queue is dependent upon three factors. Jobs with the smallest I-O and CPU time and central memory requested will be given a priority higher than normal batch (1000). The aging rate priority is added on an hourly basis. Within the input queue are batch jobs for all using organizations. Therefore, the ASD computer center has imposed a software Organizational Scheduler between the input queue and the NOS/ BE Scheduler.

Each using organization of the Cyber is authorized a percentage share of the Cyber resources. This resource use is calculated in computer resource units (CRUs) for each organization's batch and interactive

12
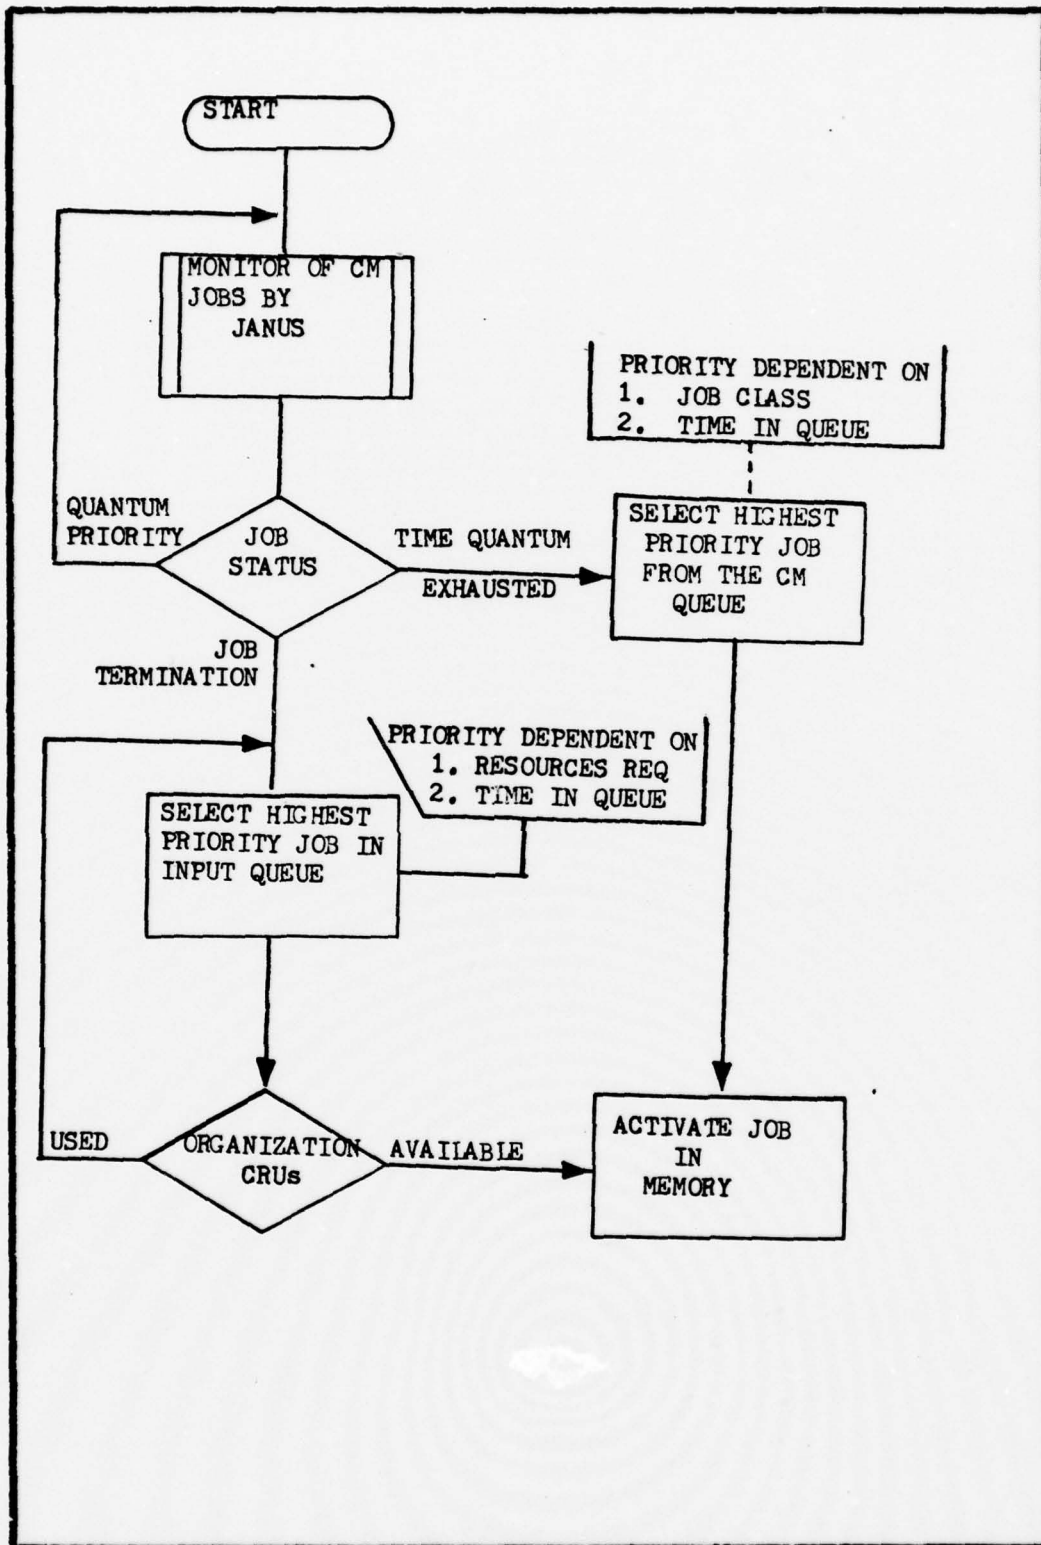
Figure 3. Job Scheduling Logic

13

jobs. CRUs are calculated for each job as a percentage of CPU plus I-0 time plus central memory used. To be selected for job activation, a job in the input queue must have the highest priority and the CRUs used by its organization must be within the CRUs authorized for this hour. The Organizational Scheduler exhibits control only over the input queue (batch jobs), even though CRU's used are a measure of batch and interactive. An organization using intercom extensively can therefore penalize batch users. These controls, such as priority job classes and aging, are used in analyzing performance.

Cyber Configuration

At the initiation point of this investigation, the Cyber 74 was being utilized in support of batch plus interactive processing only. Due to problems encountered with obtaining CLARA tapes, only a limited amount of data had been obtained. The effort was therefore expanded to analyze the Cyber as it was used in two configuration modes. The first configuration will be titled "Batch plus Interactive" and the second as "Batch Only". Each of these configurations will be explained in the chapter. The workload and performance analysis of Chapter IV shows results for each of these configurations.

Batch Plus Interactive. In this mode of operation, the ASD CDC 6600 and Cyber 74 supported the following users:

The principal users of the Cyber 74 computer system:

-Air Force Flight Dynamics Laboratory (AFFDL)

-Air Force Avionics Laboratory (AFAL)

-Air Force Institute of Technology (AFIT)

14

-Aerospace Medical Research Laboratory (AMRL)

-Air Force Human Reliability Laboratory (AFHRL)

The principal users of the CDC 6600 Computer system:

-Aerospace System Division/Other (ASD)

-Aerospace System Division/EN (ASD/EN)

-Aerospace System Division/XR (ASD/XR)

-Air Force Materials Laboratory (AFML)

-Air Force Wright Aeronautical Laboratory (AFWAL)

-Air Force Contractor Support (X)

Both the CDC 6600 and Cyber 74 operate 24 hours daily, seven days a week except for preventive maintenance. Each system supports dial-up interactive users, graphics, and remote batch terminals. Intercom use hours generally are between 0800-2200 hours, Monday through Friday and 0900-1600 hours on Saturday. These hours are continually modified in an attempt to provide better service.

Table III reflects the allocated computer resource units (CRUs) and disk space authorized each user. This CRU percentage is the control for the Organizational Scheduler previously mentioned in the discussion of job flow. The ASD computer center is a centrally-funded organization, therefore users transfer no funds to it. It is possible for an organization to exceed its CRU percentage; the only consequence is a reduction in organizational performance. The CRU is also used to compute a dollar value charge for each user. Since the systems are in effect free service, CRUs have a minimal control effect.

Additional controls placed on the Cyber system dictate that during the 0800-1600 hours time frame, all interactive jobs, and only those

15

TABLE III

ASD Using Organizations

ORGANIZATION	ORGANIZATION IDENTIFIER	CYBER 74 % CRUs	CDC 6600 % CRUs
AFFDL	D	49	1
AFAL	V	21	1
AFIT	T	12	1
AMRL	L	9	1
AFHRL	H	3	1
ASD	A	1.5	35
ASD/EN	E	.5	22
ASD/XR	R	.5	13
AFML	M	1	12
AFAPL	P	1	12
AFWAL	W	.5	1
CONTRACTOR SUPPORT	X	.5	.5

16

batch jobs requesting less than 170K octal words of memory, 400 seconds of CPU time and 1,000 seconds of I/O time, will be processed. The above was the configuration of the ASD CDC systems prior to 11 November, 1978. The replacement, also the current configuration, is now discussed.

Batch Only. Under this configuration, the CDC 6600 was designated as primarily an interactive system and the Cyber 74 became primarily a batch system although the Cyber 74 continues to service three Cybergraphic terminals until January 1979. The separation was made possible by the development of access to permanent files from either system. The stated computer center goals were to:

-Allow more time for running jobs requiring large central memory use.

-Tune each system more easily for running its own type of jobs more efficiently (i.e., fewer memory swaps for large batch jobs and less swap time-wait-time for interactive jobs).

The CRUs authorized per user remained the same, although the authorization was averaged between the two machines, which insured that the batch users of an organization will no longer be penalized if its organization's intercom users consume more than 50 percent of the CRU allocation. The redefinition of job classes for batch processing scheduling has not been defined because of the need for further evaluation. The evaluation of the Cyber workload before and after the reconfiguration is a part of this investigative effort.

III. CPEM Methodology

This chapter presents the methods used to collect workload and performance values. Then, techniques for constructing workload models and deriving system performance from this data are discussed.

Data Collection: CLARA

The data collection for workload and performance measures was made available by the Boeing Computer Service program, CLARA (Ref 9). The description of the NOS/BE accounting data and the CLARA program will be presented in the following paragraphs. Also, problems encountered with CLARA data and the reduction of that data will be examined.

Description. The CLARA program accepts data from the NOS/BE operating system Dayfile tapes, which contain all event messages generated by the operating system during normal processing. A new Dayfile tape is started every 24 hours (0001-2400). Each Dayfile message begins with the computer clock time of day followed by messages of the following three types (Ref 9:5):

- Identification data, such as job name, receipt data, job origin (batch, interactive), and account number (Ref 9:A-5).

- Initiation and termination times, such as the time a job entered the input queue, and the times that a job entered and left a control point (Ref 9: A-23).

- Quantities of computer resources used such as CPU, PPU, and I-O time, and central memory, tapes drives, and disk sectors used (Ref 9: A-25).

The CLARA program is then used to process each day's Dayfile tape to reduce its large volume down to a more managable size. These data are stored on magnetic tapes, titled Permanent Data Tapes (PDT), by CLARA.

18

For each day, the following files are produced on the PDT:

-The Summary file which is a general description of the day's activities for the system as a whole.

-The Tape-Reel file for magnetic tape activities.

-The Queue file for the times a job entered and departed the job queues.

-The Execution file which contains the job identification and computer resource usage fields.

Reference 9 gives a complete description of each of these files. The files of interest were the Queue and Execution file. In attempting to access information for each of these for a reduced data file, several problems were encountered. These will be discussed in the next section.

Problems. The initial problem with the CLARA PDT was discovered during an attempt to read the Queue and Execution files. Each daily PDT is kept by the ASD computer center for approximately two months, but the center does not use this data for any purpose. Attempts to access the files revealed that the PDT Execution file was missing for each day due to an error in the job control cards used in executing the CLARA program. This data source was therefore invalidated.

Compounding this problem was a previous decision by the ASD computer center to quit producing PDT tapes because of their nonuse. The center had assumed that the historical PDT tapes were valid and would be sufficient for workload analysis. Because the historical PDTs were invalid, the computer center agreed to resume production of PDTs until sufficient data was gathered. Due to the time delay caused by this problem, workload models are based on two weeks data. The data used will be discussed further in Chapter IV.

19

Once a week of CLARA data was produced, an analysis of the
Queue and Execution files was conducted to insure the integrity of
the formats and descriptions of each parameter in reference 9.
Table IV is a description of parameters that were deficient in the
Queue and Execution files. The methods used to obtain replacements
(if possible) for these parameters are discussed in the next section,
which is an outline of a data reduction program written to extract
those parameters needed from each day's PDT.

Data Reduction. Because the PDT consists of four files containing
143 parameters (Ref 9: Appendix A), a data reduction program was writ-
ten. The main functions performed by this program are listed below:

-For each day on the PDT, select those Queue and Execution
 records within the time period requested.

-Read and store from the Queue file for each job the time it
 entered the input queue (applicable to batch jobs only).

-Skip all overhead jobs in the Execution file, since these jobs
 are not considered as workload in this investigation.

-The account number parameter, missing for interactive jobs, was
 moved from the user name field to its correct location.

-The calculation of central memory used was derived as shown in
 Figure 4.

-The reduced job data was formatted as shown in Table V and
 written to a permanent file.

The permanent file shown in Table V was further processed to select
specific parameters for constructing workload models.

Workload Characterization

As mentioned in previous chapters, one way to define workload
is as the amount of hardware resources consumed by each job. Therefore,

20

TABLE IV

Deficient CLARA Job Parameters

Description	Action Taken
Time job printing complete	Excluded from investigation
User account number	Reference data reduction section
Central memory used	Reference data reduction section
Extended Core Storage Used	Excluded from investigation
Number of disk accesses or requests	Excluded from investigation
Intercom connect time	Excluded from investigation
Disk drive occupancy time	Excluded from investigation
Number of cards read	Excluded from investigation
Number of lines printed	Excluded from investigation
Total number of swapouts	Excluded from investigation
Swap out time	Excluded from investigation
Time spent in output queue	Excluded from investigation

The Algorithm for computation of a job's CRUs is as follows:

Number of CRUs. = a(CPU) + b(IO) + ((CPU+IO)CM)

where CPU = job's central processor time in seconds

IO = jobs's tape channel and disk access time in seconds

CM = number of central memory words required

a = CPU percentage cost ratio

b = IO percentage cost ratio

c = CM percentage cost ratio

Contained within the Execution file for each job are values

for the above CRU components of:

1. CPU sec

2. IO sec

3. (CP + IO) CM

Then CM used = $\frac{(CPU + IO) \; CM}{CPU + IO}$

Figure 4. Derivation of Central Memory Parameter

TABLE V

Reduced Job Parameters

Name	Description
Job Name	Uniquely identifies each job.
Receipt Date	The date the job entered the input queue.
Job Origin	Identifies whether the job was initiated as batch, remote batch or intercom.
Organization Identifier	Identifies the organization which submitted the job.
In-Queue	Time in seconds the job entered the input queue.
In-CP	Time in seconds the job left the input queue and entered a control point.
Out-CP	Time in seconds the job left a control point and entered the output queue.
CPU	Central Processing time in seconds used by a job.
I-O	Input-Output time in seconds used by a job.
CM	Maximum central memory in kilo-words that a job occupied, same as CMR for batch jobs.
CMR	Maximum central memory requested by a batch job.
CRUs	The number of computer resource units (Ref Figure 3) used by a job.
PPUs	The PPU time in seconds used by a job, includes I-O channel plus I-O overhead time.
Tape-Drives	Maximum number of tape drives used at one time by a program.
Disk	Maximum number of disk sectors used at one time by a program.

it was necessary to select resource consumption parameters that best represent the workload. Using SPSS (Ref 21: Chapter 18), available parameters were analyzed for their interdependence, as shown in Table VI. Using this analysis to support already existing knowledge (Ref 9, 12), the following parameters were excluded:

-The CRU parameter because it is a measure of CPU, I-O, and CM usage.

-The PPU parameter because of its dependence on I-O time.

-The Tape-Drives parameter because of the infrequent use of tapes during the 0800-1700 period.

This left CPU, I-O, CM, and Disk usage as the remaining workload characterization parameters. A remaining problem to be solved was how to model the workload using these parameters. The first model used descriptive statistics.

Descriptive Statistics. Kelly (Ref 20) defines CPE as the statistical analysis of computer performance. Descriptive statistics, such as pictorial displays of data, measures of location, and measures of variability are the most used statistical methods for CPE. However, there are problems in characterizing workload with descriptive statistics. As an example, the mean value is, without doubt, the most common measure of central tendency; but it may be the least meaningful. It is greatly affected by extreme values and may give an unrepresentative picture of typical behavior. In CPE, it is usually "typical behavior" that one cares about, and, by definition, the most typical value is the median, because it is not affected by extreme values (Ref 20: 3). To overcome the shortcomings of descriptive statistics, another workload model was developed. This model describes the workload as

24

TABLE VI

Job Parameter Correlation Analysis

	CPU	I-O	CM	CRUs	PPUs	TPDRIVE	DISK
CPU	1.00	.66	.19	.79	.31	.06	.62
I-O	.66	1.00	.25	.97	.65	.15	.75
CM	.19	.25	1.00	.29	.25	.12	.26
CRUs	.79	.97	.29	1.00	.60	.14	.77
PPUs	.31	.65	.25	.60	1.00	.28	.30
TAPE-DRIVES	.06	.15	.12	.14	.28	1.00	.16
DISK	.62	.75	.26	.77	.30	.16	1.00

groups of similar jobs or job classes using a technique called clustering.
Once the job classes are determined, the presentation of each job class
is done with descriptive statistics. The differentiation between the
two uses of descriptive statistics lies in the extra processing step
that produces job classes. The clustering technique is now discussed.

Clustering Application. The purpose of a cluster analysis is
to group observations into m subclasses, such that the differences
between members of the same class are minimized. An observation is
the job as represented by the resources it consumed. The cluster analy-
sis technique used regards each job as a point in four dimensional (CPU,
I-O, CM, Disk) Euclidean space. The jobs are assigned to a group (clus-
ter) such that the sum of the squares of the distances of a point to an
existing cluster is minimized. The methods to accomplish clustering are
dependent upon the clustering algorithm used.

The first constraint on selecting an algorithm was availability.
Dubes (Ref 15) has developed or modified eight of the most common clus-
tering techniques. Since the Air Force Avionics Laboratory had a tape
containing these programs (Ref 18), it only remained to choose one of
the eight available. An examination of the programs revealed that all
but one limited the number of observations (jobs), because all observa-
tions were required in memory at once. Because the number of jobs to
cluster numbered in the thousands, the algorithm capable of handling
this was used. This algorithm, titled WISH, is a version of Wishart's
variant on the K-means method (Ref 2: 169-170).

WISH uses a multipass approach, which starts with a set of initial
cluster centers, and assigns subsequent job observations to these centers,

26

based on a minimum distance between the observation and the cluster centers. Observations too far away from any previously defined cluster center are used to start a new cluster center. The file of observations will be reprocessed, creating and deleting clusters until the process converges. Convergence occurs when a pass is made over the data and no observations change their cluster assignment. An understanding of the functions of WISH will be better understood by examining the input to WISH. The basic input data are the observations consisting of the CPU, I-O, CM, and Disk usage parameters for each job. The clustering progess is dependent on the control parameters input to the WISH program. Discussed below are the critical input variables. Figure 5 is an accompanying flow diagram of the discussion.

Normalization. By setting this variable to 1, the observations will be normalized prior to clustering. Each of the workload parameters will be scaled by making the mean of each parameter class zero and the variance one. Changing the origin has no effect on the clustering but normalizing the spreads of the individual workload parameters assigns equal weights to all parameters and tends to transform the cloud of observations into a hypersphere (Ref 15: 21). Because the magnitude of the parameters such as disk and CM usage vary by a ratio of 50 to 1, the magnitude of the disk parameter would weight its importance too strongly in the clustering process. In order that each parameter have equal weights, the normalization control was used.

Initial Cluster Centers. The user may input the initial cluster center if he has prior knowledge of how the data might cluster. Since the possible job classes were unknown, an option in WISH which partitions

27

START

READ CONTROL PARAMETERS

CREATE INITIAL CLUSTER CENTERS

B

READ AN OBSERVATION

REWIND FILE AND READ AGAIN

END OF FILE — YES → ANY CLUSTER MOVEMENT — YES

NO

KC= CLOSER CENTER TO OBSERVATION

REMOVE CLUSTERS LESS THAN MINSIZE

NO

$DIST=$ NORMALIZED DISTANCE FROM KC TO OBSERVATION

STOP

DIST > MAX-DIST — YES → START NEW CENTER WITH OBSERVATION → B

NO

DIST < MIN-DIST — YES → UPDATE KC AND LOSING KC WITH OBSERVATION → B

TEMPORARILY IGNORE OBSERVATION → B

Figure 5. Clustering Program Logic(Ref 15: 250)

28

the observation space into 2 ** (Number of Dimensions, 4 in this case)
cluster centers was used. Later experimentation revealed that the final
clusters produced were not dependent on the initial cluster centers. How-
ever, the number of clustering iterations could be reduced with good
guesses for initial centers.

Minimum Distance. This parameter is used in adding an observation
to an existing cluster center. First, the determination of which clus-
ter an observation is closer to is made. The observation must then be
within "Minimum Distance" units or it will not be assigned to that clus-
ter.

Maximum Distance. If the observation distance is greater than
"Maximum Distance" units away, the observation will start a new cluster
center. If the observation distance is greater than "Minimum Distance"
and less than "Maximum Distance", it will be temporatily ignored.

Minimum Size. Another parameter contributing to cluster forma-
tion is the minimum number of observations that a cluster must contain
to exist as a cluster. If a cluster has less observation than "Minimum
Size" it will be discarded until the last iteration.

These last three parameters were the critical controls used in
forming clusters. By varying "Minimum Size", extreme observations
could be ignored. By specifying the "Minimum Distance" and "Maximum
Distance" the tightness or within cluster variance could be controlled.
These parameters prevented blind application of the clustering program.
Prior to using the program, the following goals were established for
cluster acceptance:

-Classify the job types with less than ten clusters.

29

-Insure that each cluster variance is a mininum for the first goal.

-In conjunction with the first two goals, maximize the number of jobs assigned to each cluster.

In attempting to group the workload into job classes, the clustering algorithm was applied twelve times before a satisfactory classification was obtained. The clustering results will be presented later.

Because WISH only prints cluster identification information, WISH was modified to add the cluster assignment number to each observation. This provided for further workload and performance analysis for each job class using programmatic methods.

General Performance Measures

Once the workload was defined, the calculation of performance values was accomplished. The performance measures of this investigation are descriptors of system efficiency and user time delays. Performance measures concerning user time delays are as follows:

-Turnaround Time is the difference between the time a job entered the input queue and the time it left a control point (batch jobs only).

-Input Queue Time is the difference between the time a job entered the input queue and the time it entered a control point (batch jobs only).

-Control Point Time is the difference between the time a job enters and leaves a control point (batch and interactive jobs).

Performance measures for system efficiency (does not include down-time) are:

-Job Throughput, which is the number of jobs executed per hour (excludes overhead jobs).

-CRU Throughput, which is the number of CRUs used per hour (excludes overhead jobs).

30

-CPU Utilization, which is the amount of CPU time consumed during the time period (0800-1700), divided by the shift time (excludes overhead jobs).

The calculation of values for the entire system, and individual organization workloads was accomplished programmatically. The application of the data reduction program, workload characterization techniques, and performance calculations is flow charted in Figure 6. The results of this application are now presented.
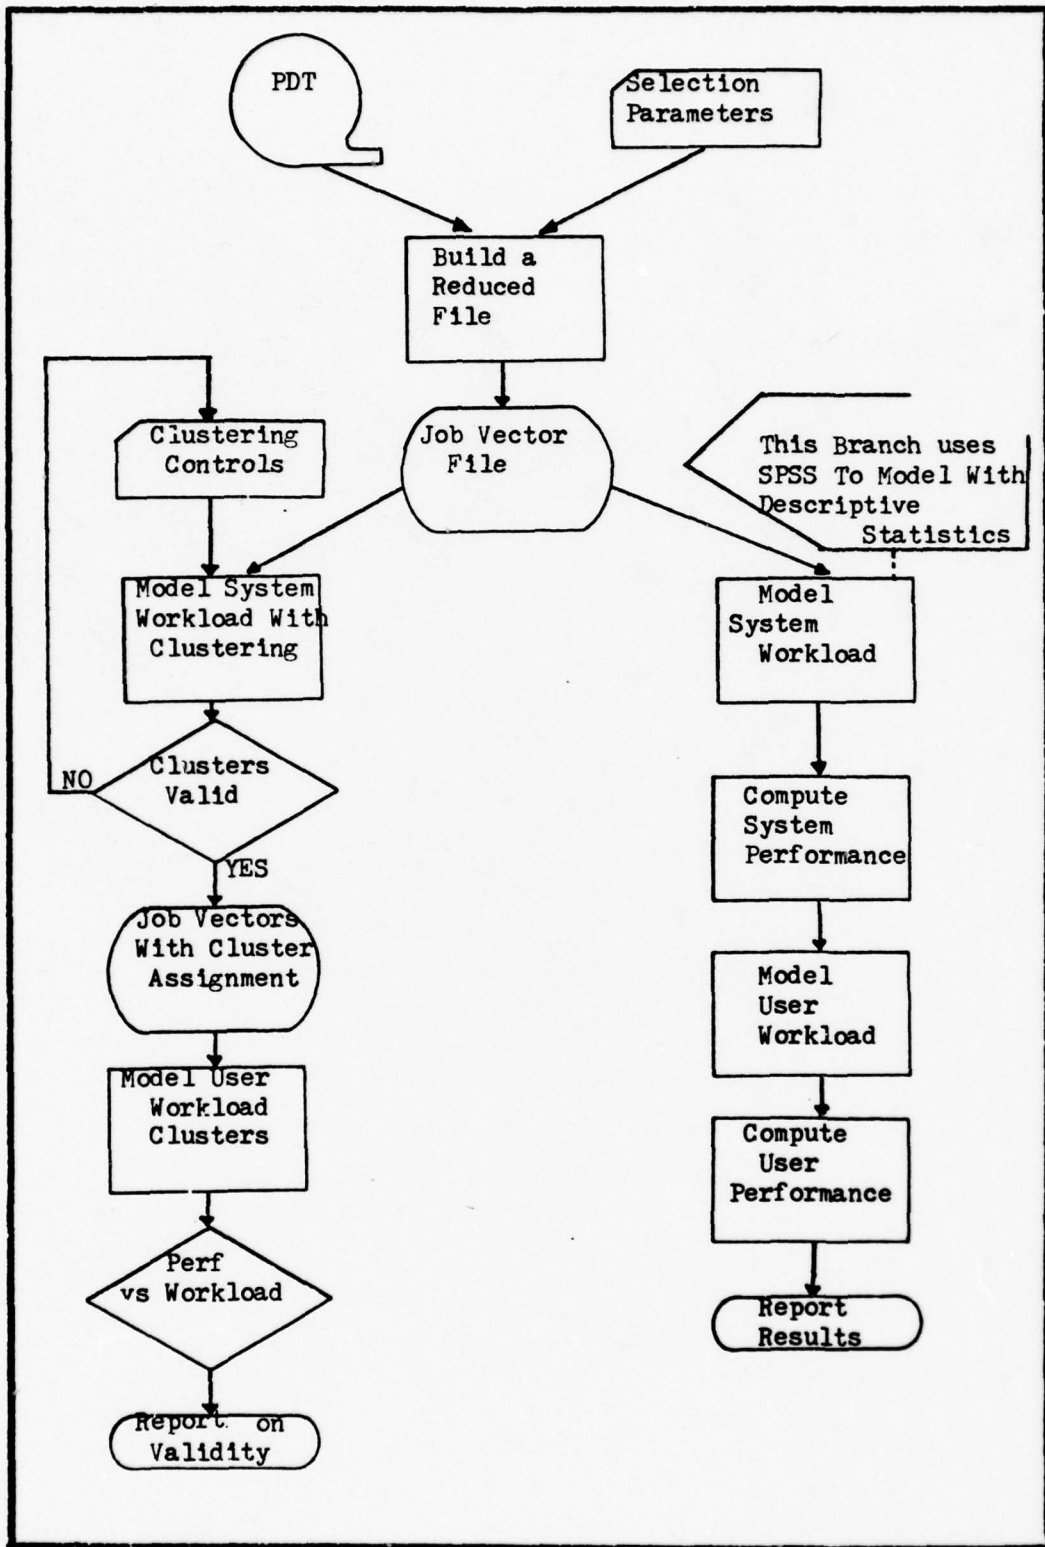
Figure 6. Methodology

32

IV. System Evaluation

This chapter presents results obtained by applying the methods of the previous chapter. CIARA data was used that was collected for a two week period during the 0800-1700 hours time frame. Data from 28 August thru 1 September, 1978 was analyzed to determine Cyber workload and performance when it was being operated in a "Batch Plus Interactive" processing mode. To analyze the Cyber in a "Batch Only" mode, data from 11 thru 15 September was used. This chapter will show results for the system as a whole (all organizations), followed by the workload and performance for each individual organization (user).

System Workload

Because the primary objective of this research was to characterize individual user workloads and performance, the question of why investigate the total system workload might arise? The answer is that the magnitude of any individual user workload is relative to the total workload! Therefore, this section will present the workload and performance of the system, saving the individual organization results for later. As previously mentioned, modeling of the workload was done from two approaches, descriptive statistics, and cluster analysis. Each of these results will be shown for the system as operated in the two processing modes.

Table VII, supported by Figures 7a. - 7d., presents statistics for CPU, I-O, and disk quantities consumed from 28 Aug-1 Sep. Analysis of this data indicates that distributions of CPU, I-O, and disk usage are positively skewed with the means of each in regions of low density.

33

TABLE VII

System Workload Statistics For 4669 Jobs (28 Aug- 1 Sep)

PARAMETER	CPU (sec)	I-0 (sec)	CM (K words)	DISK (sectors)
MINIMUM	.1	.1	.1	0
MAXIMUM	1178.0	3500.0	57.0	54000
MEAN	16.7	31.4	10.8	500
.95 CONFIDENCE INTERVAL	15.6 to 17.7	28.6 to 34.2	10.6 to 11.1	447 to 549
STANDARD DEVIATION	37.1	97.1	8.6	1700
MEDIAN	4.4	10.0	9.7	180
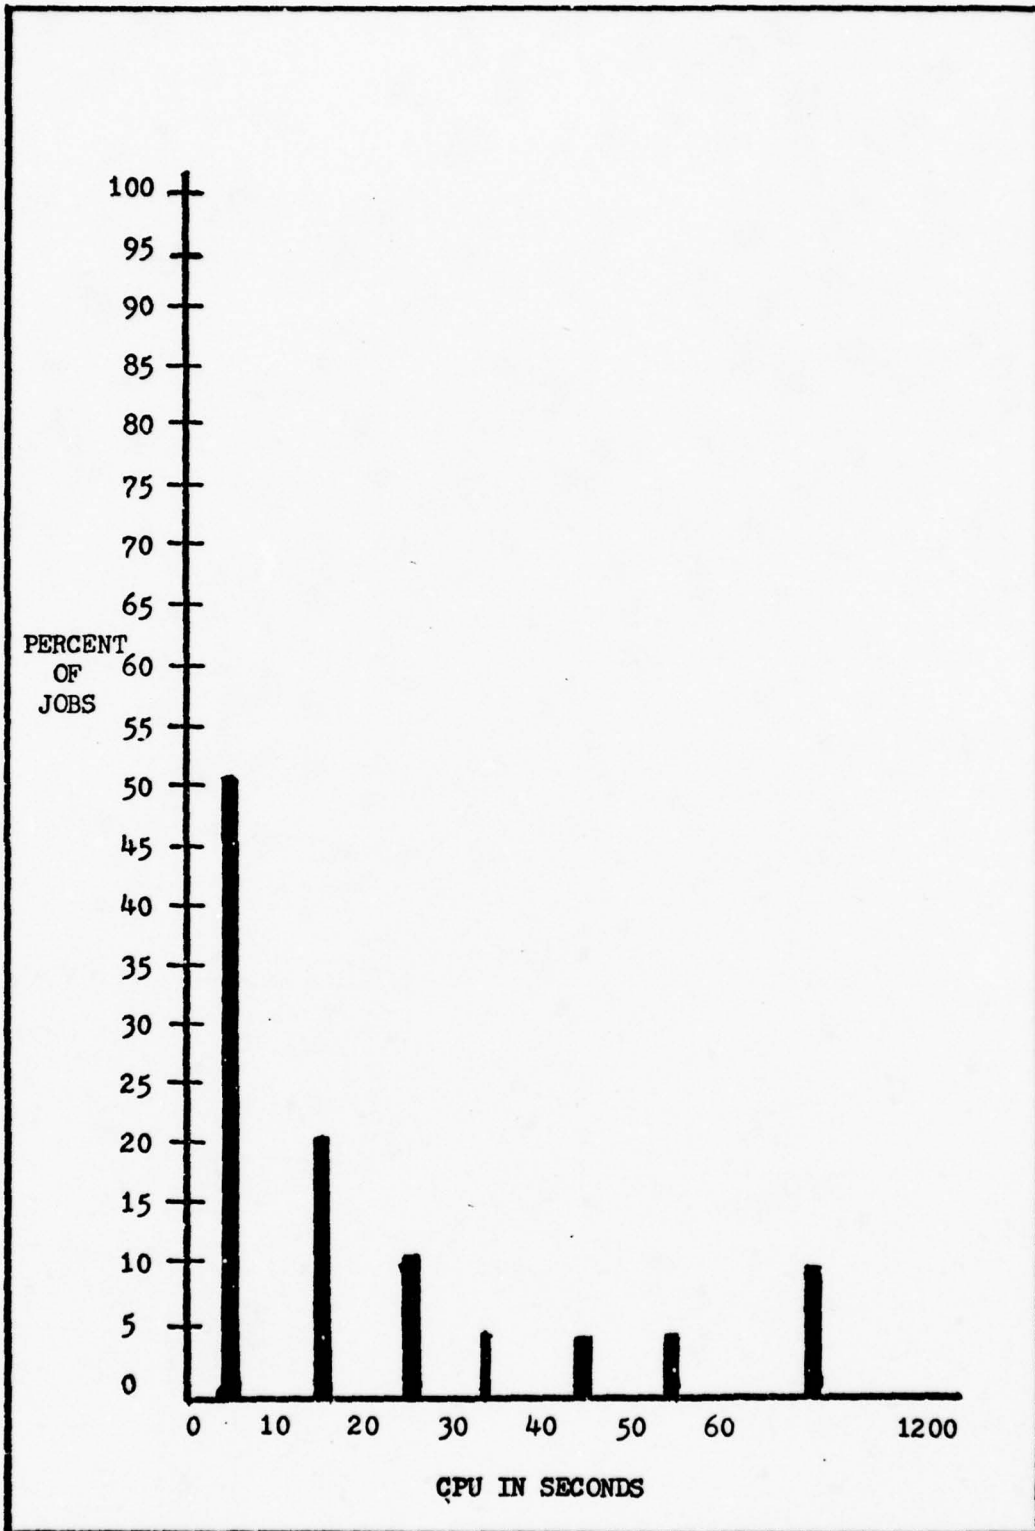
Figure 7a. Resource Usage Distribution - CPU (28 Aug - 1 Sep)

35
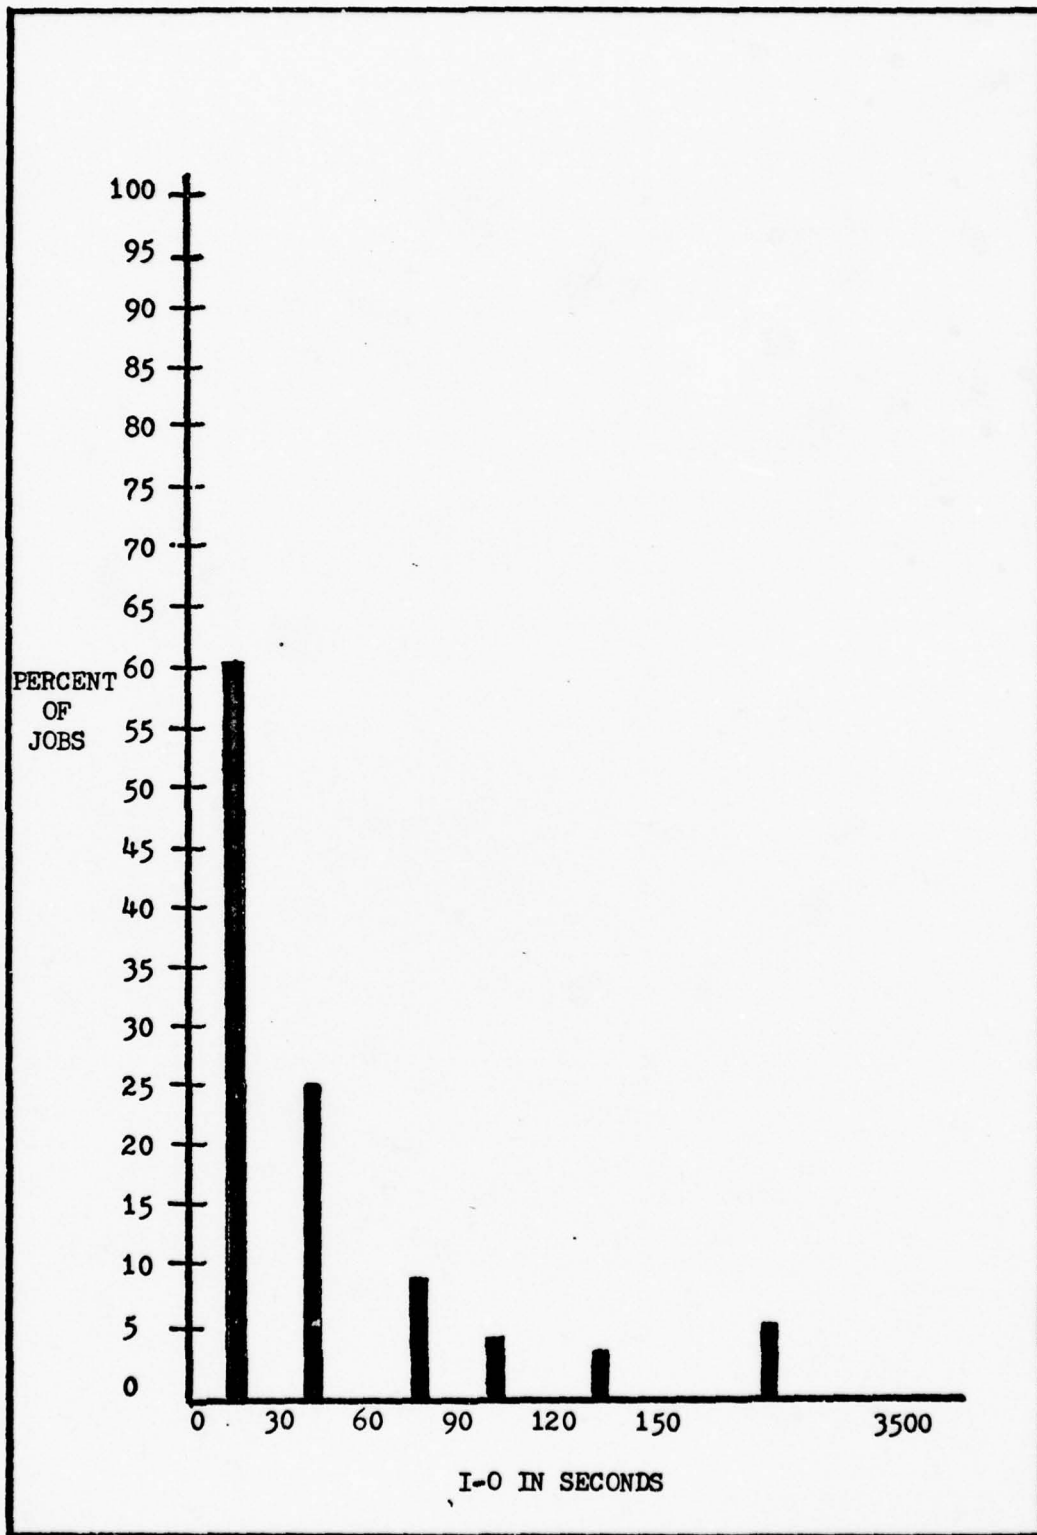
Figure 7 b. Resource Usage Distribution - I-O (28 Aug - 1 Sep)

36

Figure 7 c. Resource Usage Distribution - CM (28 Aug-1 Sep)

37

Figure 7d. Resource Usage Distribution - Disk (28 Aug - 1 Sep)

38

TABLE VIII

System Cluster Descriptions - Resource Means (28 Aug- 1 Sep)

CLUSTER	JOBS IN CLUSTER	CPU (sec)	I-O (sec)	CM (K words)	DISK (sectors)
1	1358	2	5	3	40
2	1078	3	9	15	300
3	264	8	28	17	550
4	184	27	31	6	300
5	311	76	20	11	450
6	340	16	25	25	1000
7	485	11	15	5	200
8	319	82	170	16	2500
9	330	14	125	17	1000

TABLE IX

System Cluster Descriptions - Usage (28 Aug - 1 Sep)

CLUSTER	JOBS IN CLUSTER	% BATCH	CPU	I-O	CM	DISK
1	1358	21	LOW	LOW	LOW	LOW
2	1078	94	LOW	LOW	MEDIUM	LOW
3	264	72	LOW	MEDIUM	MEDIUM	MEDIUM
4	184	7	MEDIUM	MEDIUM	LOW	LOW
5	311	29	HIGH	LOW	MEDIUM	MEDIUM
6	340	97	MEDIUM	MEDIUM	HIGH	HIGH
7	485	8	MEDIUM	LOW	LOW	LOW
8	319	34	HIGH	HIGH	MEDIUM	MEDIUM
9	330	49	MEDIUM	HIGH	MEDIUM	HIGH

Also, the CM distribution exhibits no central tendencies, suggesting that means do not reflect typical usage. The median, which shows that a majority of jobs consume small resource amounts, is a more apt descriptor. Looking at these statistics does not give a complete feel for what the Cyber workload was. Because the Cyber is a job processor, a better view of the workload would consider the types of job classes executed. A way to do this is proposed by Agrawala (Ref 1), who advocates the use of clustering for workload visualization, and as a method for generating probabalistic workload values, for input to performance simulation models (Ref 17). In this investigation, the job clusters are used only to portray the workload.

Tables VIII and IX are two different views of the same cluster descriptions (job classes). The means and cluster assignments are a product of the clustering program. The percentage of batch jobs was calculated after the clustering was complete. Determination of whether the mean was low, medium or high was based on the medians in Table VII. These tables define the job class workloads, but a few observations are pertinent. Jobs in clusters 1 and 2 comprise more than 50% of the jobs executed, with the jobs in each one utilizing small resource amounts. This agrees with the observations made above. However, the rationale for job class resources can be obtained from the cluster means. As an example, the greater CM usage by cluster 2 over 1 is due to cluster two's higher percentage of batch jobs, which generally request more memory than interactive jobs. Each of the remaining clusters is unique in terms of the resources consumed. Jobs in clusters 8 and 9 are larger

41

jobs, with those in 8 being CPU bound and those in 9 I-O bound. The stability of job clusters on a daily basis is verified by Table X. These results indicated that a workload characterization of one day may be as valid as one week, since no daily trends were seen. Prior to calculating the performance for the 28 Aug-1 Sep period, the workload for 11-15 Sep was determined.

The Cyber primarily executed batch jobs during the 11-15 Sep time frame, due to the changed operating mode effective 11 Sep. Table XI and Figures 8a.-8d. show that the consumption of memory, I-O, and disk usage increased with a slight drop in CPU usage. This tends to agree with the ASD computer center's goal of running larger jobs. By looking further at the cluster description of the workload in Tables XII and XIII, it can be seen that the number of jobs using larger resource amounts increased. The type of jobs that decreased during the 0800-1700 shift, were those using medium resource amounts, which were the interactively spawned jobs. Once again, Table XIV points out the consistency of job types executed daily. The workload imposed on the system is of interest, but the key concern is what performance the system delivers under this workload.

System Performance

The performance delivered by a system can only be rated good or bad if an objective measure is available. Because no quantitative performance standards are set by the ASD computer center, the determination of performance quality is limited to what seems reasonable. However, two quantitative comparisons were available. First, the difference in

42

TABLE X

System Cluster Description - Daily Cluster Distribution (28 Aug-1 Sep)

CLUSTER	WEEK	MON	TUES	WEDS	THURS	FRI
1	29	31	27	31	28	28
2	23	22	26	21	22	23
3	6	6	5	5	6	5
4	4	3	4	4	4	4
5	7	7	7	9	5	5
6	7	8	7	6	6	9
7	10	10	10	10	12	10
8	7	6	6	7	8	7
9	7	6	6	8	8	7

TABLE XI

System Workload Statistics For 4501 Jobs (11-15 Sep)

| PARAMETER	CPU (sec)	I-O (sec)	CM (K words)	DISK (sectors)
MINIMUM	.1	.1	.1	0
MAXIMUM	1231.3	3453.2	75.0	111440
MEAN	17.9	38.9	15.1	950
.95 CONFIDENCE INTERVAL	16.6 to 19.2	35.8 to 41.9	14.8 to 15.3	827 to 1068
STANDARD DEVIATION	43.6	103.8	9.4	4118
MEDIAN	4.1	12.5	14.7	300
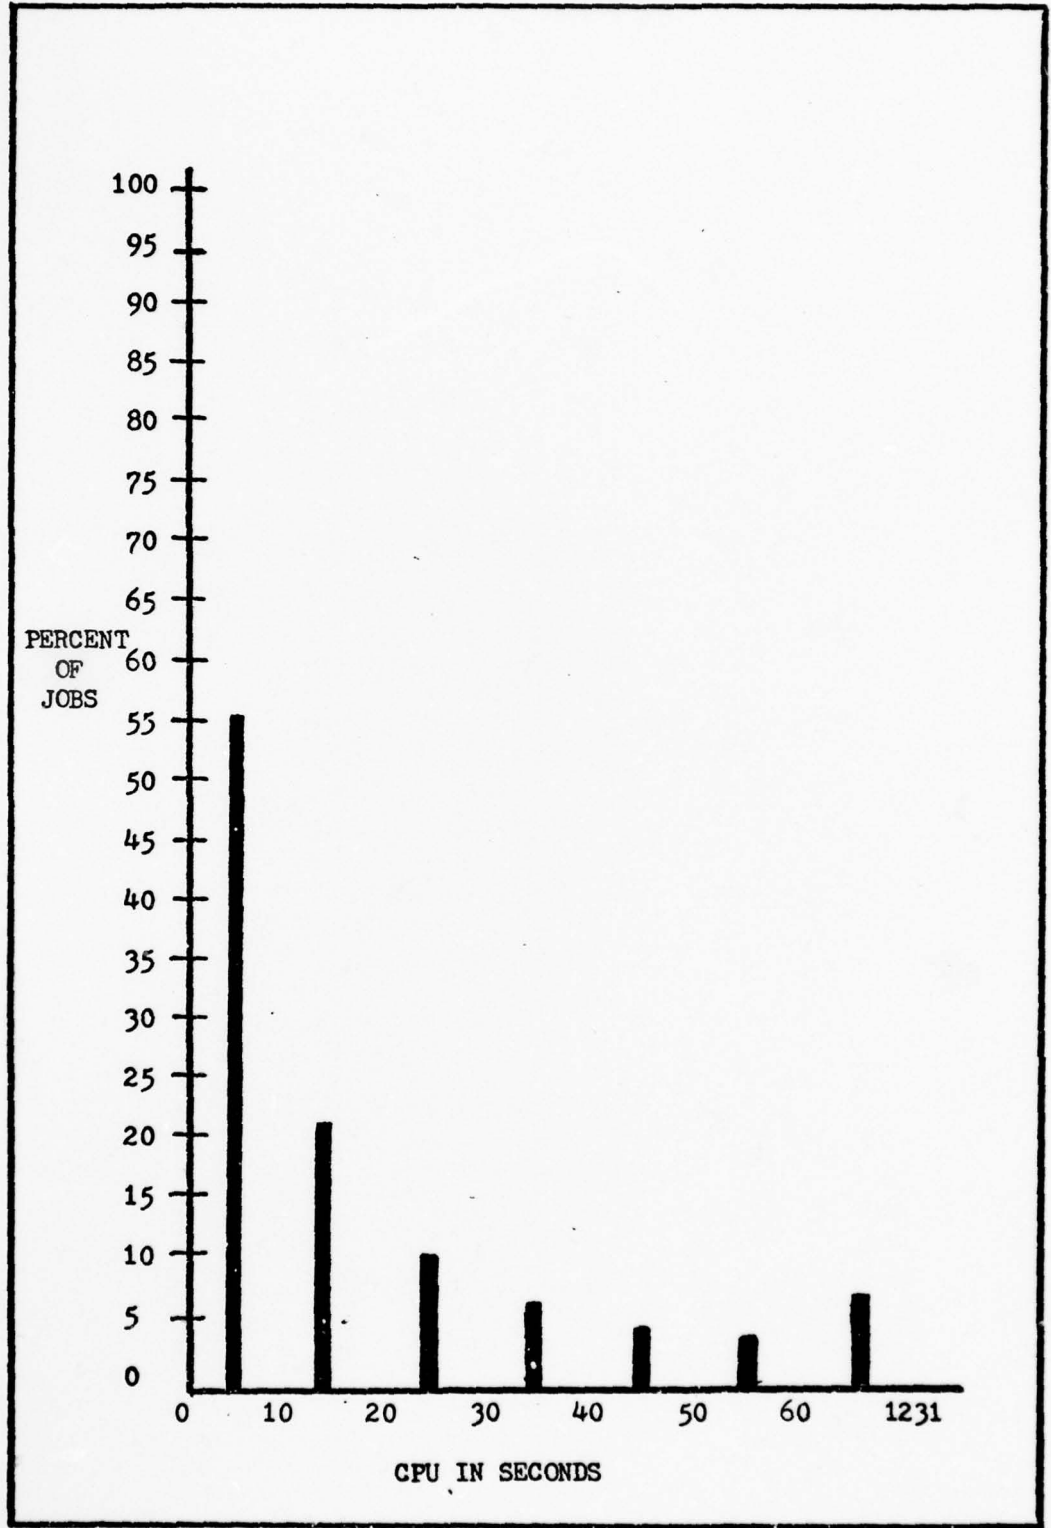
Figure 8a. Resource Usage Distribution - CPU (11-15 Sep)
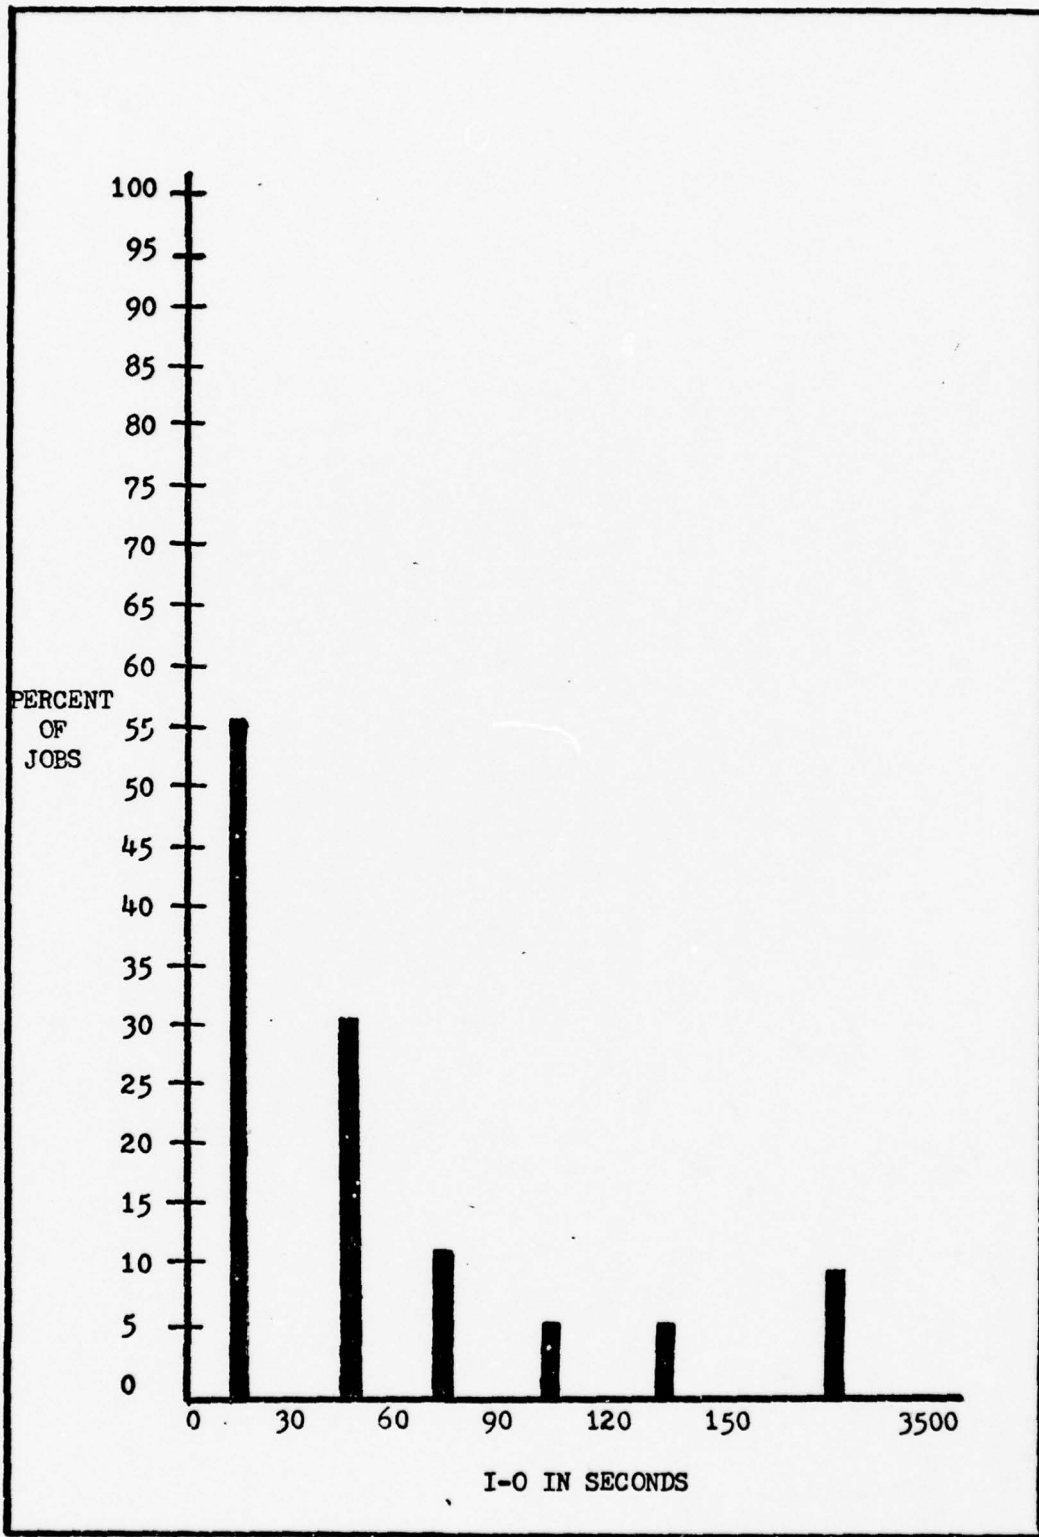
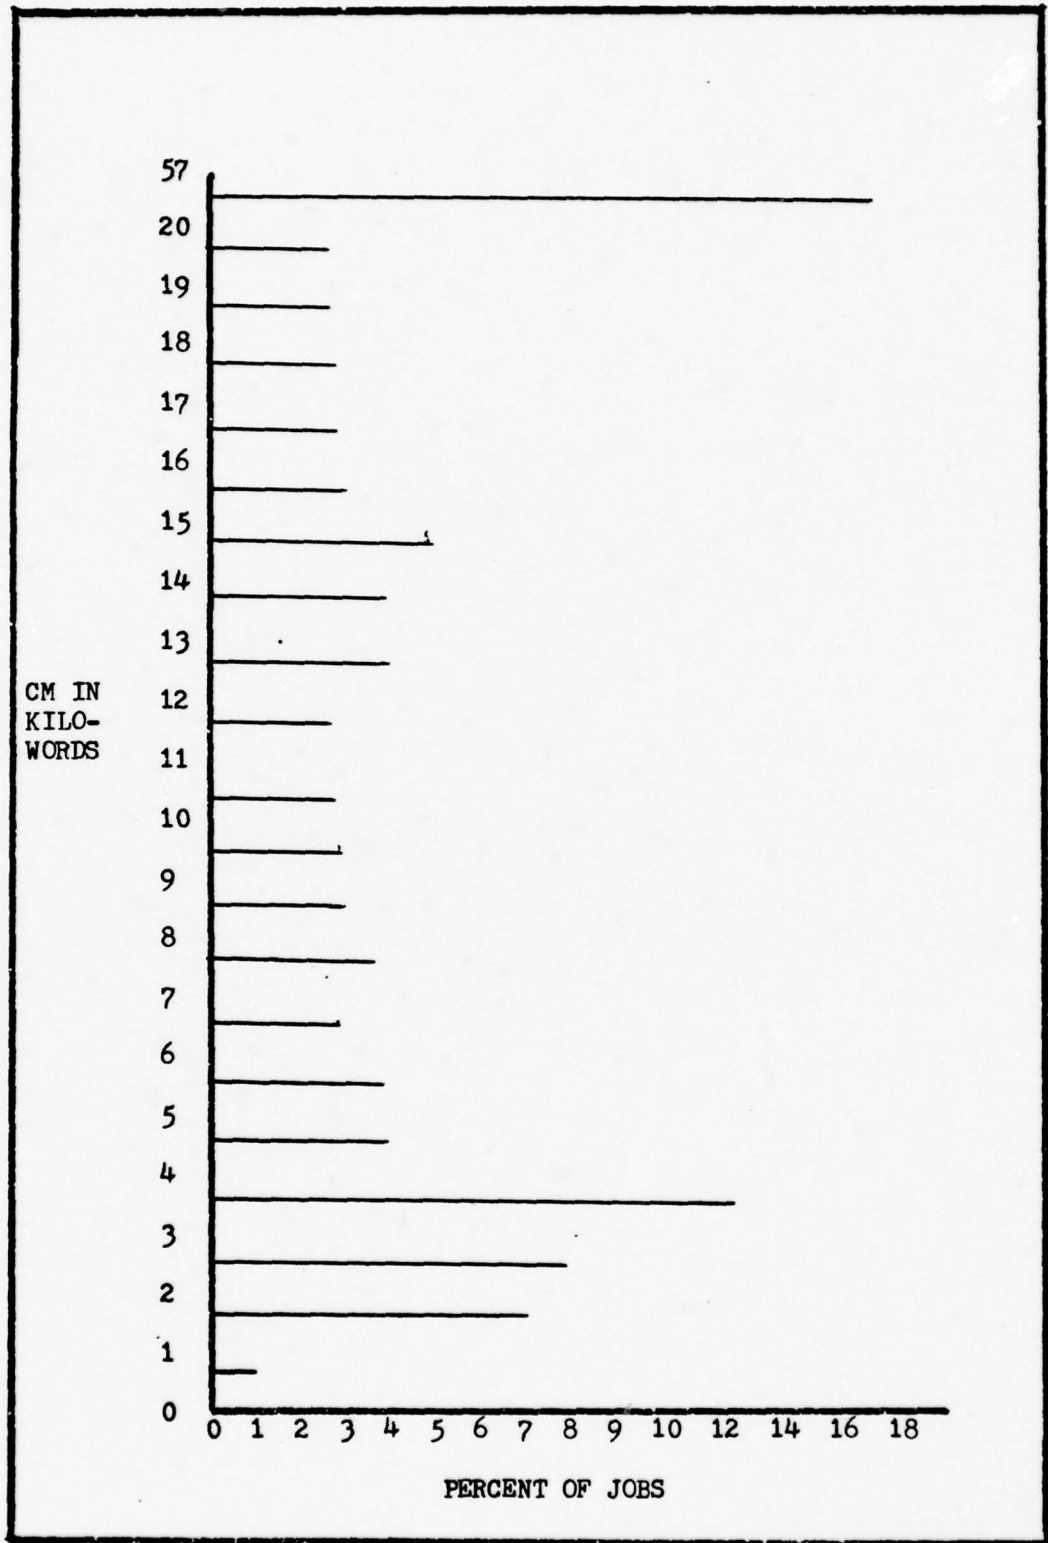Figure 8b. Resource Usage Distribution - I-O (11-15 Sep)

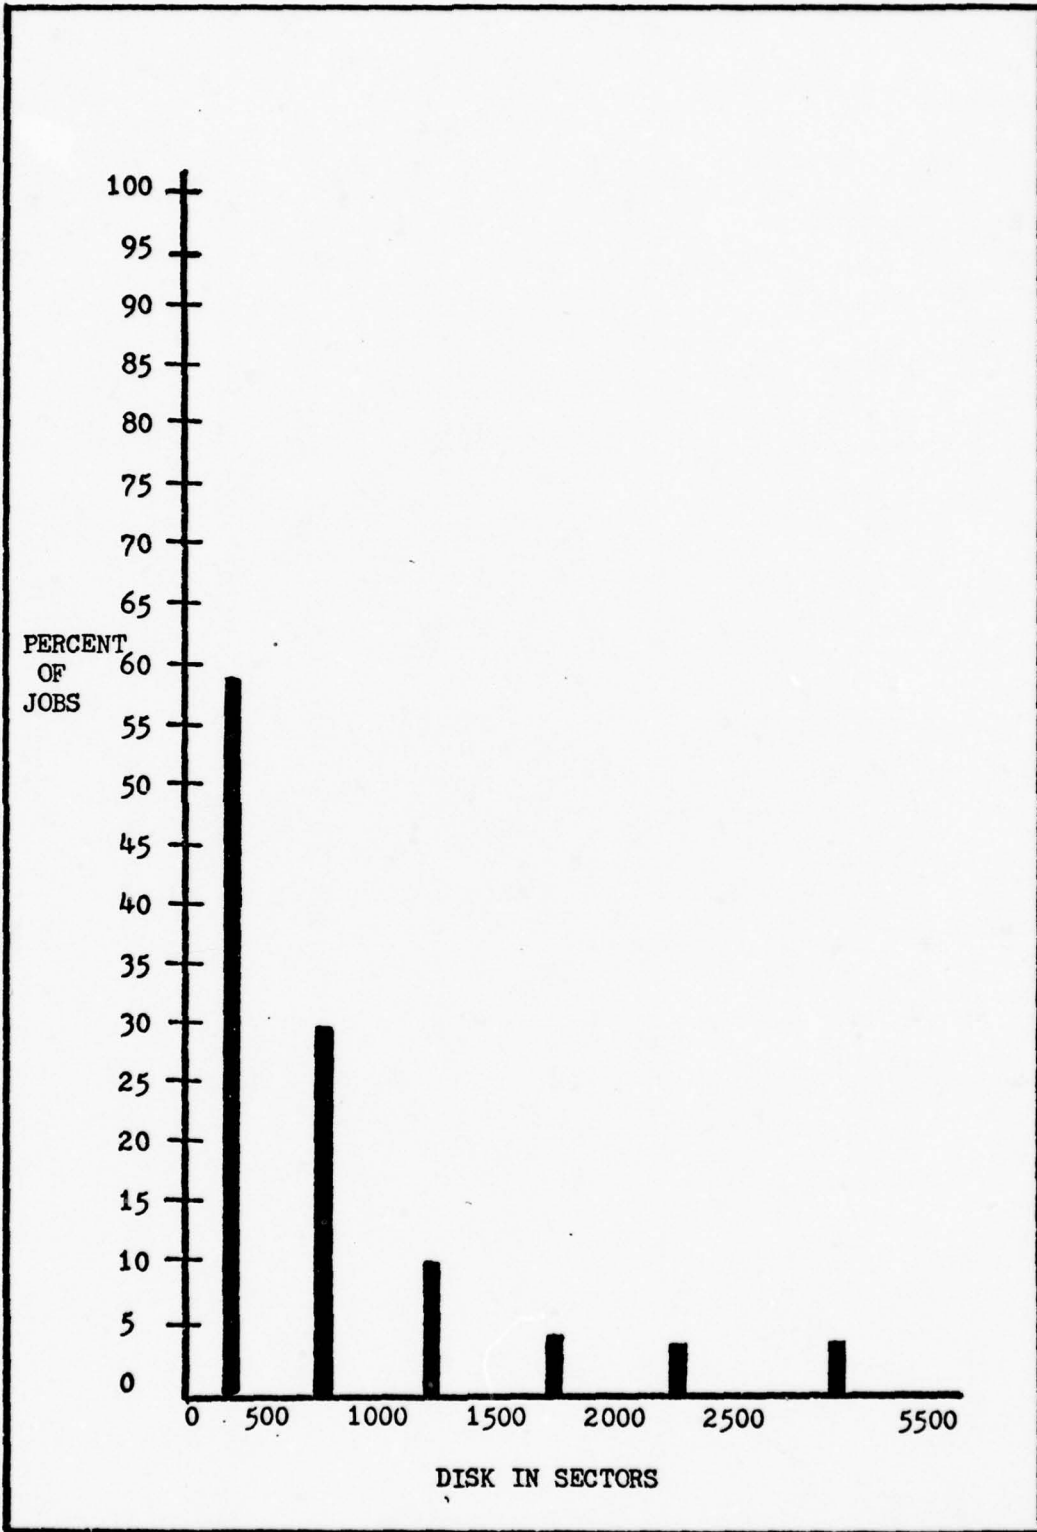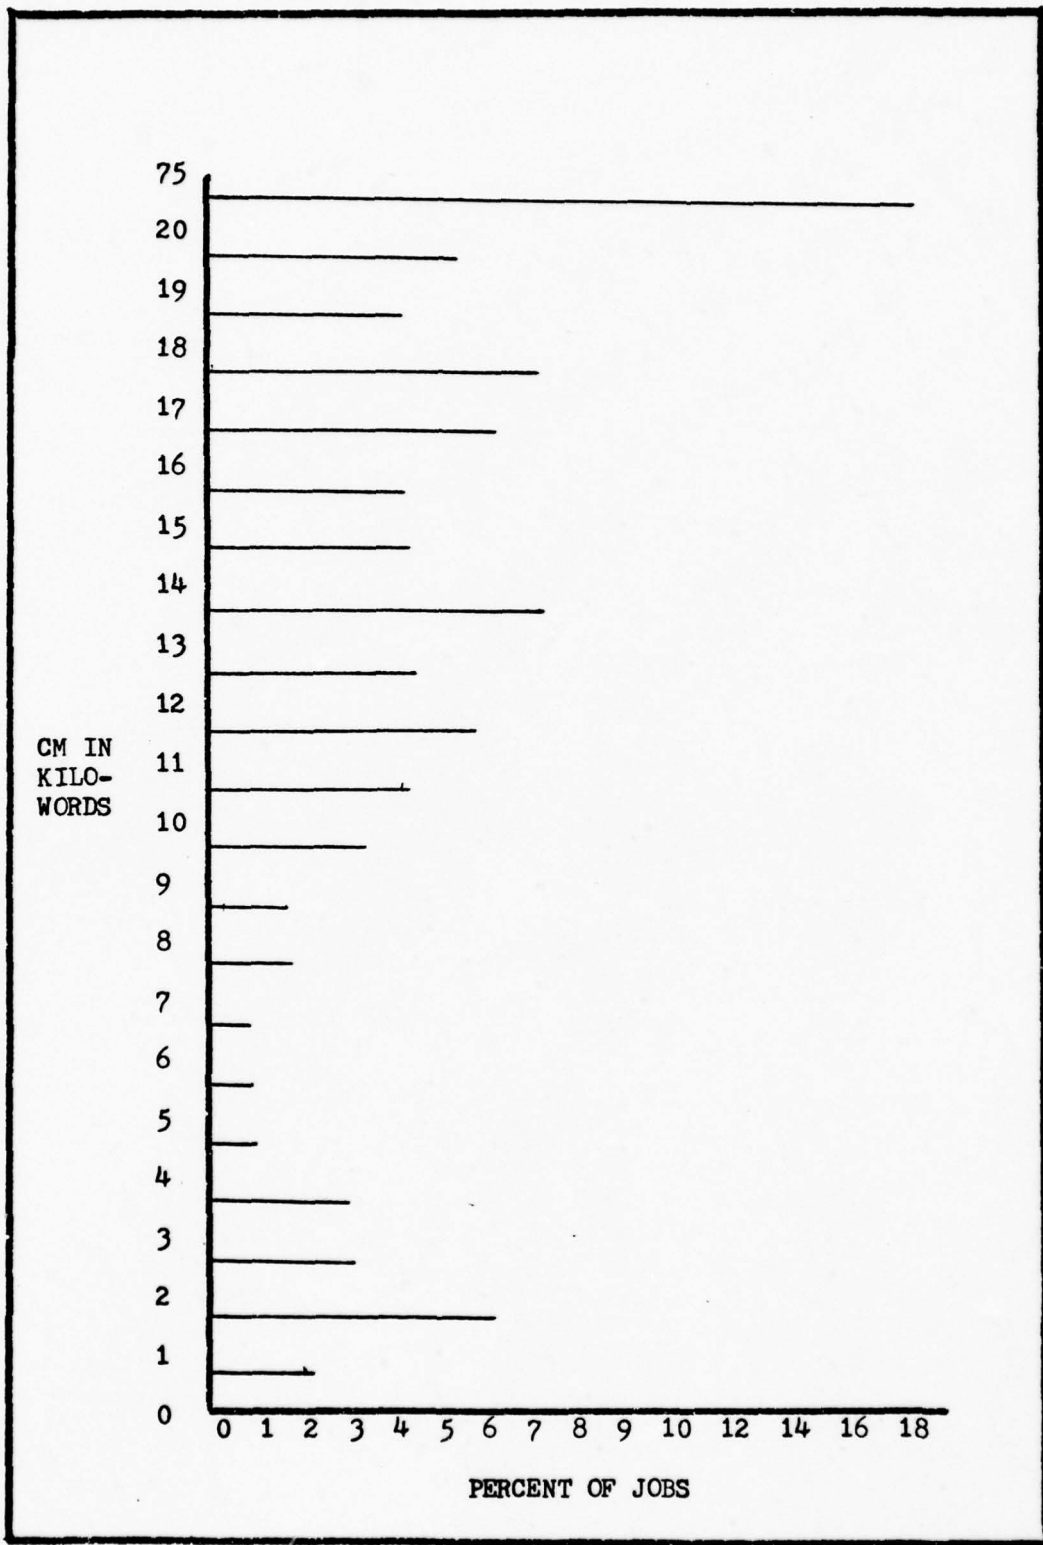Figure 8c. Resource Usage Distribution - CM (11-15 Sep)

47
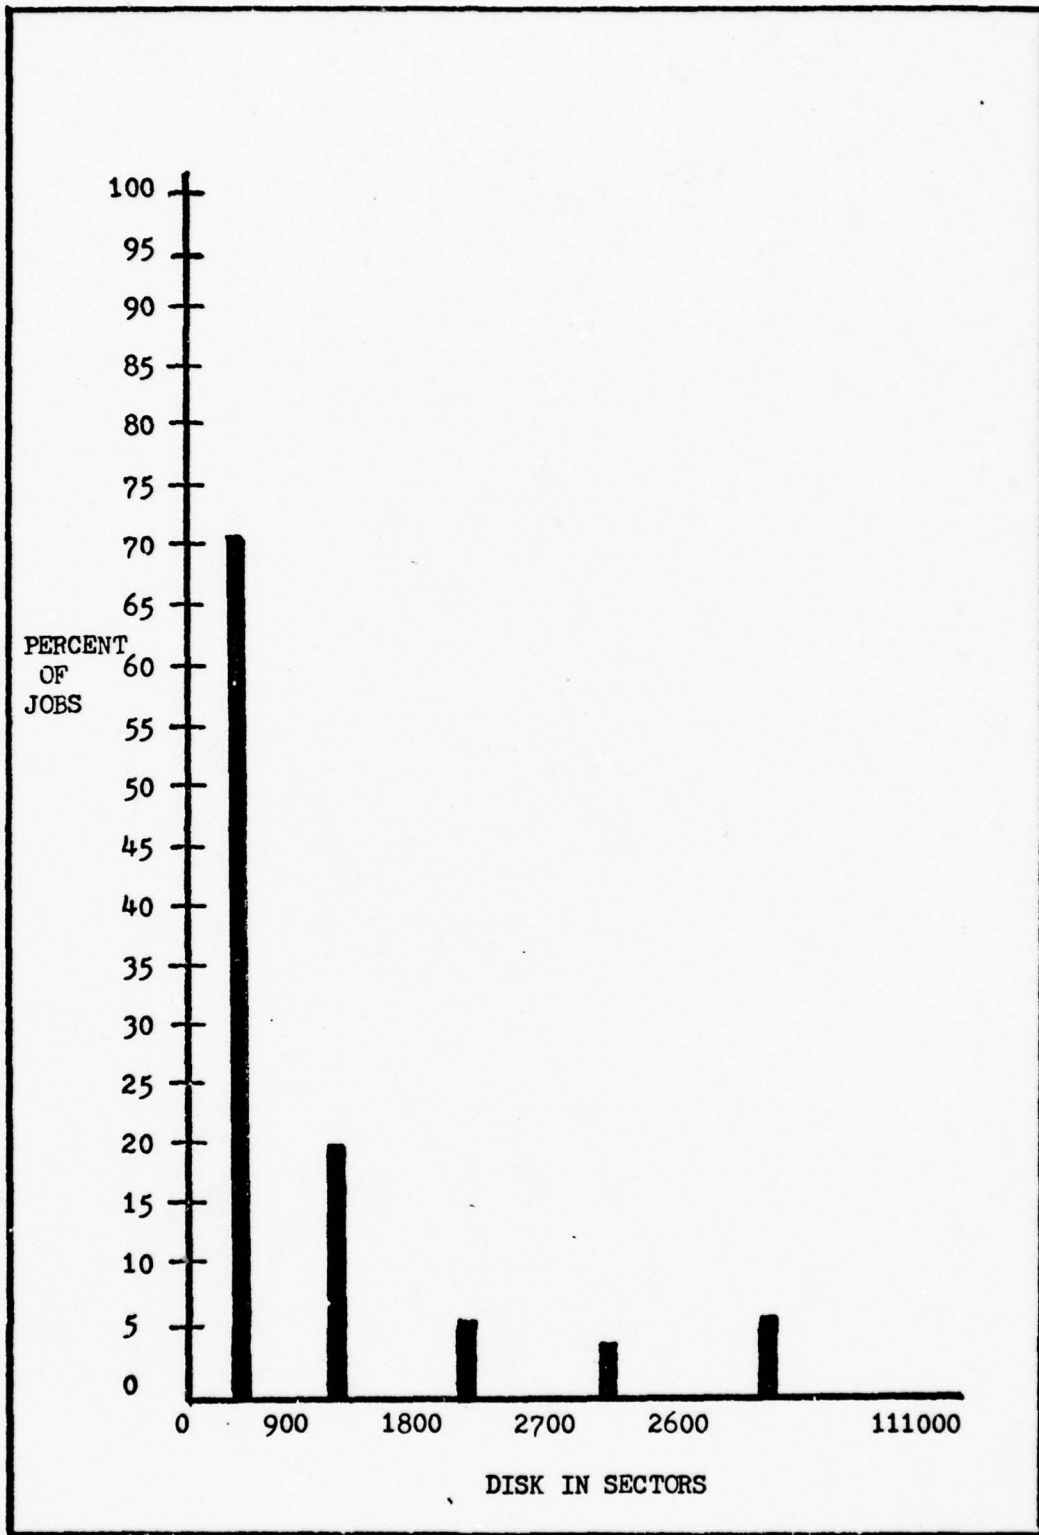
Figure 8d. Resource Usage Distribution - Disk (11-15 Sep)

48

TABLE XII

System Cluster Descriptions - Resource Means (11-15 Sep)

CLUSTER	JOBS	CPU (sec)	I-O (sec)	CM (K words)	DISK (sector)
1	1361	3	9	12	300
2	753	.5	5	2	90
3	702	6	14	20	400
4	459	12	30	19	900
5	386	55	238	20	3000
6	292	84	25	16	900
7	310	19	80	19	3500
8	238	62	48	35	2000

TABLE XIII

System Cluster Descriptions - Usage (11-15 Sep)

CLUSTER	JOBS	% BATCH	CPU	I-O	CM	DISK
1	1361	99	LOW	LOW	LOW	LOW
2	753	85	VERY LOW	VERY LOW	VERY LOW	LOW
3	702	100	LOW	LOW	MEDIUM	LOW
4	459	99	MEDIUM	MEDIUM	MEDIUM	MEDIUM
5	386	98	HIGH	HIGH	MEDIUM	HIGH
6	292	89	HIGH	LOW	MEDIUM	MEDIUM
7	310	97	MEDIUM	HIGH	MEDIUM	HIGH
8	238	100	HIGH	MEDIUM	HIGH	HIGH

TABLE XIV

System Cluster Description - Daily Cluster Distribution (11-15 Sep)

CLUSTER	WEEK	MON	TUES	WEDS	THURS	FRI
1	30	31	31	32	29	30
2	17	15	18	16	17	16
3	16	21	18	17	10	15
4	10	7	6	9	15	10
5	9	7	8	7	11	9
6	6	5	6	7	7	9
7	7	8	7	7	6	5
8	5	7	6	6	4	5

performance between the two Cyber operating modes could be compared. Discussed later are differences in performance among the individual using organizations.

Table XV shows an improvement in all performance measures, except throughput in jobs/hour and CPU utilization, for 11-15 Sep. This result was expected because of the near elimination of the higher priority interactive jobs. The control-point occupancy time shows a significant decrease, indicating that the number of swaps between memory and the CM queue is near zero. This was accomplished without a drastic drop in CPU utilization which might have been expected if a few I-O jobs had monopolized memory. This mix of jobs is borne out by analyzing the increase in small resource jobs that accompanied the increase in large resource jobs. Looking at the performance figures of Table XV does not indicate which job types benefitted or suffered due to the changed operation. An attempt to derive performance measures for job classes for 28 Aug-1 Sep is now discussed.

Table XVI contains turnaround statistics for each job class. An examination of the means and standard deviation indicates a great deal of variability in the turnaround for each cluster. Table XXVI in Appendix A supports this picture of unpredictable performance for a given job class by showing the frequency distribution of turnaround time for the jobs in each cluster. An effort to relate median turnaround (Table XVI) to resources used (Table IX) for each cluster was also unsuccessful. The net result is that turnaround is not related to the derived job classes. This is not unlikely, because performance is dependent not only on the workload, but also on the system design. The design criteria

52

TABLE XV

System Performance Means

	28 Aug - 1 Sep	11-15 Sep
TURNAROUND (MIN) (BATCH JOBS)	60.6%	43.4%
INPUT-QUEUE (MIN) (BATCH JOBS)	53.2	38.6
CONTROL-POINT (MIN) (ALL USER JOBS)	18	4.6
THROUGHPUT-JOBS/HR (ALL USER JOBS)	117	102
THROUGHPUT-CRUs/HR (ALL USER JOBS)	3440	3800
CPU UTILIZATION/SHIFT (ALL USER JOBS)	54.2	50.8

TABLE XVI

System Turnaround (Min) By Cluster (28 Aug- 1 Sep)

CLUSTER	MAXIMUM	MEAN	STANDARD DEVIATION	MEDIAN	SKEWNESS
1	923	51	95	11	4.9
2	1100	46	95	7	5.5
3	397	72	91	32	1.8
4	683	87	184	24	3.2
5	351	76	78	48	1.8
6	1165	81	138	30	4.3
7	187	34	45	14	1.8
8	868	124	137	71	2.3
9	440	63	82	28	2.4

discussed in Chapter II are restated for an understanding of job

turnaround on the Cyber (Ref Figure 3, Chapter 2).

-The priority of a job affects its turnaround; batch is lowest,
intercom is higher, and graphics is the highest. This priority
is critical for initiation, swapping, and execution.

-The resources requested (memory, disk, tapes) can delay the
initiation of a batch job; lower resources requested get
quicker entry.

-The two criteria above are constrained by availability of an
organization's CRU's; the highest priority batch job will not
be initiated if the alloted CRUs for his organization have been
exceeded.

Because the 28 Aug-1 Sep workload was a mixture of interactive and

batch jobs, all of the above criteria impacted performance. To sup-

port or refute the conclusion that performance is not related to job

classes, the 11-15 Sep clusters were analyzed. Since the workload for

11-15 Sep was predominantly batch, the effect of priority on perform-

ance should be minimized.

Table XVII reflects turnaround quantities for the 11-15 Sep job

classes shown in Table XIII. Because the input-queue dalay of turn-

around is affected the most by the above criteria, the input-queue

statistics, Table XVIII, and frequency distribution (Ref Appendix A,

Table XXVII), were analyzed for performance variability. Once again

the conclusion reached was that the other criteria affect performance

to the extent that a particular job class will not have a definite turn-

around time. Further analysis was conducted to determine if a job class

could expect to spend a mean or median time in a control point. Tables

XIII and XIX are used for the following discussion. Prior to comparing

the two tables, remember that control point time is dependent on CPU

55

TABLE XVII

System Turnaround (Min) By Cluster (11-15 Sep)

CLUSTER	MAXIMUM	MEAN	STANDARD DEVIATION	MEDIAN	SKEWNESS
1	650	33	63	7.4	3.29
2	415	26	51	7.8	3.12
3	435	30	57	6.8	3.1
4	320	44	63	14.3	1.9
5	505	58	72	26.7	2.5
6	320	44	67	17.8	2.4
7	455	58	75	25	2.0
8	490	48	69	18.4	3.0

TABLE XVIII

System Input-Queue (Min) By Cluster (11-15 Sep)

CLUSTER	MAXIMUM	MEAN	STANDARD DEVIATION	MEDIAN	SKEWNESS
1	640	32	62	6.7	3.3
2	400	25	49	6.2	3.3
3	410	29	56	5.4	3.2
4	315	41	62	1.07	1.9
5	485	38	66	10.4	3.0
6	310	38	67	11.0	2.4
7	430	51	73	18.8	2.0
8	485	39	68	10.6	3.2

TABLE XIX

System Control-Point (Min) By Cluster (11-15 Sep)

CLUSTER	MAXIMUM	MEAN	STANDARD DEVIATION	MEDIAN	SKEWNESS
1	42	1.1	2.8	.6	7.3
2	61	2.1	6.3	.2	4.9
3	34	1.3	2.4	.8	4.9
4	29	3.0	3.2	1.9	3.7
5	64	15.5	13.5	11.1	1.6
6	44	6.0	6.3	4.0	3.0
7	64	6.4	8.5	3.6	4.0
8	50	7.9	7.8	4.9	2.5

and I-O time used. A job will spend longer in a control point if
CPU time is larger because the job loses the CPU everytime its time
quantum expires. A job also loses the CPU everytime it ask for an
I-O operation. Thus high I-O and/or high CPU time will mean a longer
time in the control point. Examination of Table XIII versus XIX shows
a high correlation between the median control-point time and CPU/I-O
time. As an example, cluster 5 jobs use a high amount of both CPU
and I-O time resulting in the greatest control point time. Cluster 2
uses very low amounts, thus, its jobs spent the shortest time in a
control point. All other clusters support the contention that control-
point time is related to the resource usage (CPU, I-O) job class. This
is expected, because once a job is in a control point, scheduling cri-
teria and priority (since this data was predominately batch) have no
effect on control point time. However, the control point time is a
small portion of the total turnaround time. This leaves the mean de-
scriptor for all combined job classes as the best measure of system
turnaround. Another way to view Cyber performance is the service
received by the individual users.

Organization Workload

Table XX contains overall job percentages by organization with
AFFDL, AFAL, AFIT, and AMRL consuming the largest share of Cyber re-
sources. Remember that the CRU is the ASD computer center's resource
control, it can be seen that AFIT processed twice as many jobs as AFAL,
but consumed only 2% more CRUs. This suggests that individual jobs
processed for AFAL consume more CM, I-O, and CPU than AFIT jobs. To

59

TABLE XX

User Workload Statistics For 4669 Jobs (28 Aug- 1 Sep)

ORGANIZATION		% JOBS	USER JOB % INTERACTIVE	BATCH	PERCENT OF ALL RESOURCES USED CPU	I-O	CM	CRUs	DISK
AFFDL	*	33	60	40	36	33	31	35	40
AFAL	*	12	54	46	13	18	13	16	14
AFIT	*	27	48	52	20	19	26	18	16
AMRL	*	13	30	70	11	8	15	9	10
AFHRL	*	0	-	-	-	-	-	-	-
ASD		7	48	52	9	13	7	12	12
ASD/EN		3	80	20	4	3	2	3	2
ASD/XR		2	79	21	4	2	2	3	2
AFML		1	57	43	1	1	2	1	1
AFAPL		1	70	30	1	2	1	2	2
AFWAL		.5	.5	.5	.5	.5	.5	.5	.5
CONTRACTOR		.5	.5	.5	.5	.5	.5	.5	.5

* = Primary Cyber User

60

determine the type of jobs executed for each organization, the percent of previously derived cluster types are shown in Table XXI. The larger percentage of cluster 2 jobs (small resources used) by AFIT and the greater share of cluster type 9 jobs (larger resource used) by AFAL supports the above contention. The information in Tables XX and XXI is of interest to each user, but to limit the discussion, the remainder of this section will focus on AFIT workload and performance. The analysis used will also apply to the other user results. Of further interest is how the workload changed when the Cyber was converted to a "Batch Only" processing mode.

Because the interactive processing for each organization was transferred to the CDC 6600, the workload results of Tables XXII and XXIII are limited to each organization's batch workload. Comparing Tables XX and XXII revealed that the precent of jobs executed did not change as much as might be expected. As an example, AFIT's job percentage only dropped 3% indicating that AFIT was getting more batch jobs executed during the 0800-1700 shift as a result of the reconfiguration. Also Table XXIII points out the large number of small resource usage batch job executed by AFIT. Remaining to be determined is what performance was received by each organization and whether workload was the primary criteria in determining this performance.

Organization Performance

The following discussion will require the use of workload Tables XX and XXI in support of the performance results of Table XXIV. An examination of the turnaround results in Table XXIV show a wide disparity

61

TABLE XXI

User Cluster Percentage (28 Aug- 1 Sep)

CLUSTER NUMBER

USER	1	2	3	4	5	6	7	8	9
AFFDL	34	16	5	4	7	7	12	8	6
AFAL	27	16	7	4	6	7	9	9	16
AFIT	28	29	6	5	7	7	11	3	5
AMRL	19	40	5	3	6	11	7	5	3
AFHRL	-	-	-	-	-	-	-	-	-
ASD	27	30	5	3	4	5	8	7	11
ASD/EN	34	7	3	6	10	4	14	16	6
ASD/XR	30	6	1	2	9	1	14	30	7
AFML	33	15	5	2	5	13	7	8	13
AFAPL	23	11	9	6	8	2	11	8	23
AFWAL	25	38	6	0	0	0	0	6	25
CONTRACTOR	35	0	6	0	0	53	0	6	0

62

TABLE XXII

User Workload Statistics For 4501 Jobs (11-15 Sep)

ORGANIZATION	% JOBS	USER JOB % INTERACTIVE	BATCH	PERCENT OF ALL RESOURCES USED CPU	I-O	CM	CRUs	DISK
AFFDL	22	7	93	24	22	19	22	23
AFAL	7	0	100	8	7	6	8	5
AFIT	24	8	92	16	9	22	11	10
AMRL	10	0	100	6	6	11	6	6
AFHRL	-	-	-	-	-	-	-	-
ASD	10	0	100	15	22	11	19	26
ASD/EN	9	0	100	11	14	12	14	11
ASD/XR	3	7	93	3	5	3	4	3
AFML	6	0	100	6	3	7	4	4
AFAPL	8	1	99	11	8	9	9	8
AFWAL	-	-	-	-	-	-	-	-
CONTRACTOR	-	-	-	-	-	-	-	-

Note: No longer are there Primary Cyber Users

63

TABLE XXIII

User Cluster Percentages (11-15 Sep)

Cluster Number

USER	1	2	3	4	5	6	7	8
AFFDL	26	26	10	8	10	8	9	5
AFAL	20	30	8	10	13	4	10	5
AFIT	44	13	13	12	2	12	4	0
AMRL	37	11	24	12	4	2	3	7
AFHRL	-	-	-	-	-	-	-	-
ASD	34	9	15	9	16	1	11	4
ASD/EN	15	10	19	9	17	3	8	18
ASD/XR	27	18	13	6	21	3	8	5
AFML	12	14	41	17	0	4	5	7
AFAPL	26	17	15	9	7	10	9	7
AFWAL	-	-	-	-	-	-	-	-
CONTRACTOR	-	-	-	-	-	-	-	-

TABLE XXIV

User Performance Means for 4669 Jobs (28 Aug- 1 Sep)

ORGANIZATION	AUTHORIZED CRUs %	BATCH JOBS (MIN)		ALL USER JOBS (MIN) CONTROL-POINT
		TURNAROUND	INPUT-QUEUE	
AFFDL	49	48.1	37.1	17.1
AFAL	21	57.0	44.1	24.7
AFIT	12	70.7	67.5	16.5
AMRL	9	35.4	30.4	13.5
AFHRL	3	-	-	-
ASD	1.5	130.4	124.0	16.4
ASD/EN	.5	93.1	88.7	29.6
ASD/XR	.5	77.3	66.6	29.6
AFML	1.5	80.3	67.3	18.8
AFAPL	1.5	33.4	21.5	22.6
AFWAL	.5	20.9	9.9	15.4
CONTRACTOR	.5	4.5	2.2	3.2

65

in the turnaround recieved by each user. As an example, AFIT's average machine turnaround time is 15 to 30 minutes greater than AFFDL's, AFFAL's, and AMRL's time. Becaused AFIT jobs request smaller resource amounts, they should in fact get out of the batch queue sooner. Since turnaround is a batch job indicator, all priorities are the same, leaving the use of the Organizational Scheduler as the prime criteria affecting performance. As can be seen in Tables XX and XXIV, AFIT consumes more CRUs than authorized. Therefore, AFIT is penalized in its attempts to execute batch jobs. An analysis of how the reconfiguration affected performance was done next.

An analysis for the reconfiguration impact is applicable only to AFFDL, AFAL, AFIT, and AMRL because these organizations processed all their batch load on the Cyber before as well as after the reconfiguration. Table XXV reflects the new performance averages as well as the amount of CRUs authorized. This CRU figure is the avarage of the total CRUs authorized for the 6600 and the Cyber. Comparing Table XXV to XXIV shows an improvement in each organization's performance. However the difference among organization performance remained approximately the same as before reconfiguration. AFIT still has the smallest jobs, as evidenced by the average time in the control point, yet experienced the greatest delay. This was due to AFIT still using more than its share of CRUs while the other organizations did not. Therefore, the Organizational Scheduler penalizes AFIT batch jobs by leaving them in the input queue longer. The reduction in control-point time shows the effect of eliminating the higher priority interactive job mix. This concludes the presentation of results for this investigation.

66

TABLE XXV

User Performance Means For 4501 Jobs (11-15 Sep)

ORGANIZATION	AUTHORIZED CRUs %	BATCH JOBS (MIN)		ALL USER JOBS (MIN) CONTROL-POINT
		TURNAROUND	INPUT-QUEUE	
AFFDL	25	39.6	34.4	5.1
AFAL	11	41.1	30.4	10.8
AFIT	6.5	56.1	54.6	2.4
AMRL	5	21.4	17.4	4.1
AFHRL	2	-	-	-
ASD	18.25	55.6	49.1	6.5
ASD/EN	11.25	42.1	36.8	5.3
ASD/XR	6.75	13.33	9.47	5.2
AFML	6.5	21.0	17.8	3.3
AFAPL	6.5	47.0	43.6	3.5
AFWAL	.75	-	-	-
CONTRACTOR	.5	-	-	-

V. Conclusions and Recommendations

Conclusions

The objectives of this investigation were to characterize the workload by organization and to determine the system performance for this workload. The approach taken to accomplish this was composed of three steps. First CIARA, a source of workload and performance data for the Cyber, was analyzed and programmatically processed to produce a manageable data file. Secondly, the hardware dependent resource parameters for each job were used to model the workload, by the use of descriptive statistics and cluster analysis. Finally, performance values were derived for these workload models. The conclusions reached for this workload modeling and resulting performance are now discussed.

Workload. CIARA data tapes were acceptable sources of information; however, the areas of deficiency must be corrected to improve workload and performance evaluation. Due to initial problems in obtaining the valid CIARA PDT's, the period of investigation was limited to a week of data before and after Cyber system reconfiguration. A longer period of time, preferably several months, could have provided a possible workload trend. Once the CIARA data was reduced, the workload was characterized by two methods.

Describing the workload with descriptive statistics left an incomplete picture. By defining the job classes forming the workload, better visibility of the workload was provided. The clustering technique provided a method for accomplishing this; the only difficulty

68

in the procedure was the trial and error application of the algorithm, which was need to find an optimal number of clusters (job classes). Once the system workload was quantified, the investigation objective of determining individual organization workloads was done. For each user, their percent resource consumption and job classes were presented. The dependence of performance on workload was then analyzed.

Performance. Turnaround time statistics were calculated for each job cluster. For each cluster, the goal was to be able to relate to each job class a mean or median turnaround time. This would have provided for current analysis and future prediction of performance by job type. However, analysis revealed that turnaround time was not related to job size using the selected resource demands. Other criteria, such as priority, scheduling, and CRU authorization, affected turnaround to the extent that a simple turnaround versus job class model was not possible. The CRU Organizational Scheduler degrades performance for users underallocated in CRUs such as AFIT. Elimination of the scheduler should be considered to reduce overhead and provide better performance to all jobs requesting minimum resources. This would also allow easier tuning of the system by simplying turnaround criteria to a resource demand plus priority equals turnaround function.

Although it was not possible to determine the specific turnaround for a job class, the average performance measures for total jobs processed and individual organizations showed an improvement when the Cyber was converted to a "Batch Only" processing mode. This investigation concluded that batch performance was improved, however, the resulting impact for interactive processing on the CDC 6600 was not ascertained.

69

The final conclusion of this investigation is that performance evaluation efforts are difficult, if not impossible. Further studies are required to further this investigation.

Recommendations For Futher Study

The first task to be accomplished is to persuade the ASD computer center to resume production of the CLARA PDT. Because the CLARA program is still executed daily, the only increased expense would be the need for historical tapes. One reel of tape would suffice for a month of CLARA data for both the Cyber and 6600. The collection of data for a year, 12 tape reels, would provide a more complete workload analysis file! To improve the workload data, information on interactive response times need to be produced on the CLARA tape or obtained from the Dayfile if available. In lieu of this approach, a software monitor could be developed to collect intercom data.

This would allow workload models of the system to be characterized not only by the resources consumed, but also, by whether the request for computer service was interactive or batch. This type of model is defined by Agawala (Ref 1). This would allow the formation of a probabilistic workload, where the type of request is treated as a random variable and the resources used are treated as random vectors (Ref 1:19). This to be developed model, or the batch model constructed in this investigation, could then be used to drive a simulation model of the Cyber System. This simulation model would be constructed to produce values for the performance measures of interest, such as throughput, resource utilization, or turnaround. Hamm (Ref 17) has used clustering to drive

70

simulation models from WWMCCS installations. Hughes (Ref 16) has also used the same techniques for the Air Force Avionics Laboratory's DEC-10 system. The clustering approach produces the probability of a particular job class as well as the means of resource utilization. Resources include I-O, CPU, disk, and memory which can be employed as synthetic workload to a simulation program.

Another possible use of the reduced CIARA data would be to apply multiple regression analysis to the independent variables (I-O, CPU, disk) to determine their impact on the dependent performance variables. Certainly, there are a multitude of CPE approaches; this investigation offers one approach and some tools for further research.

Bibliography

1. Agrawala, A. K., J. M. Mohr, and R. M. Bryant. "An Approach to the Workload Characterization Problem." Computer, 9:18-32 (June 1976).

2. Anderberg, M. R. Cluster Analysis For Applications. New York: Academic Press, 1973.

3. ASD Computer Center. CDC NOS/BE User's Guide. Wright-Patterson Air Force Base, Ohio: ASD Computer Center August 1976.

4. ASD Computer Center. Introduction to ASD Computer Center. Wright-Patterson Air Force Base, Ohio: ASD Computer Center, March 1978.

5. Ball, G. H. Classification Analysis. Menlo Park, Calif: Stanford Research Institute, November, 1970.

6. Bear, J. R. Workload Characterization and Measurement of the CDC Cyber 74 Computer. Masters Thesis. Air Force Institute of Technology, Wright-Patterson Air Force Base, Ohio: March 1977. (ADA-042-184).

7. Bell, T. E., B. W. Boehm, and R. A. Watson Computer Performance Analysis: Framework and Initial Phases for a Performance Improvement Effort. Rand Technical Report R-549-1-PR. Santa Monica: Rand Corp., November 1972.

8. Blitt, William J. et al. The Advanced Logistics Systems CYBER 73: A Simulation Model. Masters Thesis. Air Force Institute of Technology: December 1975. (AD-A019841).

9. Boeing Computer Services, Inc. Computer Load and Resource Analysis (CLARA) System: Language Users Manual. Publication Number BCS-40023. Seattle, Washington: BCS, December, 1976.

10. Clayton, B. B. "An Operating System and Programming Systems For the 6600." AFIPS Conference Proceedings. 26: 41-58 (Fall Joint Computer Conference, 1964).

11. Control Data Corporation. 6000 Series Computer Systems Reference Manual. Publication Number 6010000, Minneapolis, Minn.: CDC, 1975.

12. Control Data Corporation. Scope 3.4 Reference Manual. Publication Number 60307200. Sunnyvale, California: Control Data Corporation, 1975.

13. Control Data Corporation. Scope Integrated Tuning Guide Programming System Bulletin. Publication Number 60493500. Sunnyvale, California: CDC, 1975.

14. Dodson, P. O. Control Data Corporation's CYBER 70: Procedures for Performance Evaluation. Masters Thesis. Air Force Institute of Technology, Wright-Patterson Air Force Base, Ohio: May, 1974. (AD-785129).

15. Dubes, R. C. and A. K. Jain "Clustering Techniques: The User's Dilemma." Pattern Recognition, 8: 247-260 (January 1976).

16. Grenonder. U. and R. F. Tsao, "Quantitative Methods For Evaluating Computer System Performance: A Review and Proposals." Statistical Computer Performance Evaluation, edited by Walter Freiberger. New York: Academic Press, 1972.

17. Hamm, Lt. H., "Workload Characterization Using Cluster Analysis." Briefing given at H6000 CPE Conference at the Air Force Data Systems Design Center, April 1978 and personal interview June 1978.

18. Hughes, H. D. Computer Science Department, Michigan State University. Workload Characterization For the DECSYSTEM-10, A workload study for the Air Force Avionics Laboratory. Wright-Patterson Air Force Base, Ohio, August 1977.

19. Hunt, E., et al. "Who Are the Users?-An Analysis of Computer Use In a University Computer Center." AFIPS Conference Proceedings. 38: 231-238 (Spring Joint Computer Conference, 1971).

20. Kelly, John C. "The Statistical Nature of Computer Performance Evaluation." Proceedings of CPEUG: (23-26 September, 1975).

21. Nie, Norman H., et al. Statistical Package for the Social Sciences. New York: McGraw-Hill Book Company, Inc., 1975.

22. Screenivasan, K., et al. Construction and Application of Representative Synthetic Workloads. Bedford, Mass.: Mitre Corp., August 1973.

23. Svobodova, L. Computer Performance Measurement and Evaluation Methods: Analysis and Applications. New York: American Elsevier Publishing Company, Inc., 1976.

24. Watson, R. A. Computer Performance Analysis: Applications of Accounting Data. Rand Technical Report R-573-NASA/PR. Santa Monica: Rand Corp., May 1971.

Appendix A

Performance Distribution Tables

Appendix A contains frequency distribution tables for each cluster showing the percentage of jobs that fell into the performance ranges indicated. Table XXVI presents the job frequencies for turnaround time during the week of 28 Aug - 1 Sep. Table XXVII reflects job input- queue wait time distributions for 11-15 Sep. Each of these tables is in support of determining the validity of deriving a performance value for each cluster (job class).

74

TABLE XXVI

Turnaround Time Frequency Distribution By Cluster (28 Aug.- 1 Sep)

Cluster 1

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	84	4	81
7	39	39	91	2	83
14	11	50	98	1	86
21	5	55	105	1	86
28	3	58	112	0	86
35	4	63	119	1	87
42	3	65	126	0	87
49	4	70	133	1	88
56	2	71	140	1	89
63	2	74	147	2	91
70	1	74	154	0	91
77	3	78	924	9	100

Cluster 2

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	108	1	88
9	47	47	117	1	89
18	11	58	126	0	89
27	7	65	135	1	90
28	4	59	144	1	92
37	3	72	153	1	93
54	3	75	162	1	94
63	3	78	171	1	95
72	3	81	180	0	95
81	2	83	189	1	98
90	2	84	198	1	97
99	3	87	1116	3	100

TABLE XXVI (continued)

Cluster 3

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	75	3	67
3	12	12	78	1	67
6	10	22	81	1	68
9	6	28	84	1	69
12	3	31	90	2	71
15	5	37	93	2	72
18	3	39	96	2	75
21	2	41	99	1	76
24	2	44	108	1	76
27	2	45	111	1	77
30	1	47	114	1	77
33	2	49	117	1	78
36	2	51	120	1	79
39	2	53	123	1	80
42	1	54	126	1	80
45	2	56	132	1	80
48	1	57	135	1	81
51	2	58	138	1	82
54	1	59	141	1	83
60	2	61	144	1	83
63	1	62	147	1	84
66	1	63	153	1	84
69	1	63	156	1	85
72	1	64	399	15	100

Cluster 4

Minutes	Freq Pct	Cum Pct
0	0	0
6	31	31
18	15	45
30	8	54
36	8	62
42	8	69
60	8	77
126	8	85
132	8	92
684	8	100

TABLE XXVI (continued)

Cluster 5

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	78	1	63
3	6	6	81	4	67
9	3	12	93	1	69
12	1	13	99	1	70
15	3	17	102	2	72
18	7	24	105	3	75
21	1	24	108	2	78
24	3	28	114	3	81
27	1	29	117	1	82
30	1	30	126	1	83
33	3	34	132	1	84
36	2	35	135	2	87
39	1	37	141	1	88
42	3	40	147	1	89
45	8	48	156	1	90
51	3	53	168	2	92
57	1	53	192	1	93
60	2	55	249	1	94
69	2	57	273	1	96
72	3	61	321	1	97
75	1	63	327	2	99
			354	1	100

Cluster 6

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	117	1	79
9	13	35	126	2	80
27	7	42	135	2	82
36	10	52	144	2	85
45	5	57	153	2	86
54	2	60	162	1	87
63	3	63	171	2	89
72	5	68	180	1	90
81	4	71	189	1	92
90	2	74	216	1	92
99	2	75	234	0	92
108	2	77	1120	8	100

TABLE XXVI (continued)

Cluster 7

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	42	3	68
2	13	13	48	3	71
4	18	32	52	3	74
6	3	34	54	5	79
8	5	39	70	5	84
10	5	45	74	3	87
12	3	47	76	3	89
16	3	55	104	3	92
18	3	58	136	3	95
20	3	61	146	3	97
22	3	63	188	3	100
30	3	66			

Cluster 8

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	161	3	75
7	5	5	168	1	76
14	5	9	175	2	78
21	7	17	182	2	80
28	5	21	189	1	81
35	11	32	203	1	82
42	3	35	210	1	83
49	3	38	224	1	83
56	5	42	252	1	84
63	5	47	259	2	86
70	1	48	280	2	88
77	4	51	294	2	90
84	1	52	301	1	91
91	2	54	343	1	92
98	1	55	357	1	93
105	3	58	371	1	94
112	2	60	385	1	94
119	1	61	420	1	95
126	3	63	448	1	96
133	2	65	455	1	97
140	2	67	469	1	98
147	2	29	482	1	99
154	4	72	875	1	100

78

TABLE XXVI (continued)

Cluster 9

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	100	1	82
4	9	9	108	2	84
8	5	14	112	1	84
12	10	24	116	2	85
16	9	33	128	1	87
20	2	35	132	1	88
24	7	42	136	1	88
28	6	48	144	1	89
32	4	52	152	1	89
36	1	52	164	1	92
40	4	55	168	1	92
44	2	58	184	1	93
48	1	59	200	1	93
52	2	61	212	1	94
56	4	65	232	1	94
60	1	66	236	1	85
64	2	69	264	1	96
68	1	70	280	1	96
72	2	71	308	1	97
76	1	73	348	1	98
80	2	75	364	1	98
84	2	76	380	1	99
88	1	77	384	1	99
92	1	78	440	1	100
96	2	81			

TABLE XXVII

Input-Queue Time Frequency Distribution By Cluster (11-15 Sep)

Cluster 1

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	65	1	84
5	39	39	70	1	85
10	13	52	75	1	86
15	9	61	80	0	86
20	6	67	85	0	85
25	4	71	90	0	87
30	2	73	95	1	88
35	3	75	100	1	88
40	2	78	105	1	89
45	2	89	110	1	90
50	2	81	115	1	90
55	1	82	120	1	91
60	1	83	655	9	100

Cluster 2

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	65	1	87
5	42	42	70	0	88
10	11	53	75	0	88
15	8	50	80	1	89
20	6	66	85	0	89
25	4	70	90	1	90
30	4	74	95	1	90
35	5	79	100	0	91
40	2	81	105	0	91
45	2	83	110	1	92
50	1	84	115	1	93
55	1	85	405	7	100
60	1	85			

TABLE XXVII (continued)

Cluster 3

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	60	1	85
5	44	44	65	1	86
10	11	55	70	1	87
15	7	61	75	1	88
20	5	67	80	1	88
25	5	72	85	0	88
30	3	75	90	1	90
35	2	77	95	0	90
40	3	80	100	0	90
45	2	82	105	0	91
50	1	83	110	1	91
55	1	84	440	9	100

Cluster 4

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	60	1	75
5	32	32	65	2	77
10	15	47	70	2	79
15	5	52	75	1	80
20	4	55	80	1	81
25	4	59	95	2	83
30	4	63	100	2	86
35	2	66	105	0	86
40	2	68	110	0	87
45	2	70	115	1	87
50	2	52	120	1	88
55	2	74	320	12	100

81

TABLE XXVII (continued)

Cluster 5

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	60	2	79
5	30	30	65	2	81
10	13	44	70	2	83
15	11	54	75	1	84
20	3	57	80	1	84
25	7	65	90	1	85
30	3	68	100	1	85
35	2	72	105	1	86
40	2	72	110	2	88
45	2	74	115	1	89
50	2	75	490	11	100
55	2	78			

Cluster 6

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	75	0	83
5	27	27	85	0	83
10	15	42	90	2	85
15	12	54	95	2	85
20	9	63	100	0	86
25	5	68	105	0	86
30	4	72	110	1	87
35	4	75	120	1	88
40	1	77	130	2	90
45	3	80	150	2	92
55	2	81	315	8	100
60	1	82			

TABLE XXVII (continued)

Cluster 7

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	70	1	74
5	26	26	75	2	76
10	9	35	80	3	79
15	9	44	85	0	79
20	5	49	95	1	80
25	5	54	105	0	81
30	5	59	110	0	81
35	2	61	115	1	82
40	4	65	120	1	83
45	2	67	125	0	84
50	3	70	130	1	85
55	2	71	135	2	87
60	1	72	140	2	89
65	1	74	435	11	100

Cluster 8

Minutes	Freq Pct	Cum Pct	Minutes	Freq Pct	Cum Pct
0	0	0	75	1	80
5	32	32	80	2	82
10	12	44	90	0	82
15	9	54	95	1	84
20	4	58	100	2	86
25	4	62	110	1	86
30	3	66	115	2	89
35	4	69	120	2	89
40	3	72	125	1	90
45	1	73	130	1	91
50	3	77	135	2	93
55	1	79	150	0	93
60	1	79	115	1	94
			490	6	100

Jimmy W. Thompson was born on 28 May 1946 in Fayetteville, Tennessee. He graduated from high school in Fayetteville, Tennessee in May, 1964 and enlisted in the USAF in August of that same year. Various enlisted assignments included duty as a computer systems instructor at Sheppard AFB, Texas. He entered the Airman's Education and Commissioning Program at Louisiana Technological University in September of 1971 and graduated with a Bachelor of Science in Computer Science in August 1973. Upon completion of Officer Training School in December 1973, he was commissioned a 2nd Lt. in the USAF. The next three years were spent as a computer systems analyst at the USAF Data Systems Design Center. He entered the AFIT School of Engineering in June 1977 and recieved the degree of Master of Science in Computer Systems in December 1978.

        Permanent address: Route 2
                                      Fayetteville, Tennessee 37334

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER AFIT/GCS/EE/78-20	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Computer Performance Evaluation of Individual User Workloads on the CDC Cyber System		5. TYPE OF REPORT & PERIOD COVERED
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s)		8. CONTRACT OR GRANT NUMBER(s)
9. PERFORMING ORGANIZATION NAME AND ADDRESS Air Force Institute Of Technology (AFIT-EN) Wright-Patterson AFB, Ohio 45433		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS ASD Computer Center (ADRP) Aeronautical Systems Division (AFSC) Wright-Patterson AFB, Ohio 45433		12. REPORT DATE December, 1978
		13. NUMBER OF PAGES 95
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		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

Approved for public release; IAW AFR 190-17

JOSEPH P. HIPPS, Major, USAF
Director of Information   1-23-79

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Computer Performance Evaluation   Clustering Analysis
Workload Characterization   Cyber 74
Systems Analysis
Performance Monitoring

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This paper describes computer performance evaluation and measurement techniques applied to a Control Data Corporation Cyber 74 system, located at Wright-Patterson AFB, Ohio. The performance of a computer system is influenced by its hardware configuration, system software, man-machine interaction, and application program workload. The focus of this investigation is to determine how application workload affects performance. Application workload is measured by the magnitude of job demands placed on

DD FORM 1473 EDITION OF 1 NOV 65 IS OBSOLETE
1 JAN 73

the system hardware resources; for example, CPU and I-O time, and central memory used. System accounting data provides resource usage values as well as related delay times for each job. Workload is then characterized by statistics for each job class and by total job statistics. The grouping of jobs into classes is done using clustering, a multivariate data analysis technique. Batch turnaround time is then analyzed for its dependence on workload. The methodology used and the results obtained provide performance conclusions and a background for further research.

4 — 79 END