CONFIGURATION PLANNING

ON AN ICL COMPUTER UTILIZING A

STOCHASTIC NETWORK ANALYSIS PACKAGE

by

~~IAN GRENVILLE DOUGLAS~~ KINGON

DISSERTATION

submitted in partial fulfilment of the

requirements for the degree

MASTER OF SCIENCE

in

COMPUTER SCIENCE

in the

FACULTY OF SCIENCES

at the

RAND AFRIKAANS UNIVERSITY

SUPERVISOR : DR J D ROODE

JANUARY 1982

ABSTRACT

This dissertation details the implementation of SNAP, a stochastic network analysis package, as the basis of an in-house computer configuration planning facility. The work was performed at Head Office, Gold Fields of South Africa Limited, Johannesburg, South Africa (GFSA) during the period April 1980 to December 1981.

SNAP was developed by the Institute of Applied Computer Science at the University of Stellenbosch, Stellenbosch, South Africa. The implementation of SNAP at GFSA signalled the first in-house SNAP facility, and the first SNAP implementation on an ICL computer (although implementation had been in progress at another ICL site since 1979).

Although this dissertation is very specific in nature, it is intended to provide an insight into the methodology employed in planning and implementing an in-house configuration planning facility. An overview of multiclass queueing network models and the SNAP package is provided, although no attempt is made to explain the stochastic theory of queueing networks in any detail.

Attention is thereafter focussed on the various phases of the project. Problems were encountered in monitoring performance data, and these are looked at in some depth. The question of workload characterization and the difficulties of producing a satisfactory GFSA classification strategy are then presented.

The model design, calibration and validation stages are explained using the GFSA model. Thereafter, use of the model for prediction purposes is illustrated by means of a number of examples.

Finally, the memory management model is discussed – main memory does not form part of the SNAP model and has to be dealt with as a separate issue.

SINOPSIS

In die skripsie word bespreek die implementeering van SNAP, 'n
stogastiese netwerkontledingspakkie, as basis vir 'n konfigurasie
-beplanningsfasiliteit. Die navorsing is gedoen gedurende die periode
April 1980 tot Desember 1981 te Hoofkantoor, Gold Fields van Suid
Afrika Beperk, Johannesburg, Suid Afrika (GFSA).

SNAP is ontwikkel deur die Instituut vir Toegepaste Rekenaarwetenskap
te Stellenbosch Universiteit, Stellenbosch, Suid Africa. Die SNAP
implementeering te GFSA was die eerste gebruik van SNAP as 'n
binnenshuise fasiliteit asook die eerste SNAP implementeering op 'n
ICL rekenaar (alhoewel 'n ander ICL installasie sedert 1979 besig is
met die implementeering).

Alhoewel die skripsie baie spesifiek van aard is, poog dit om insig
te gee in die metodologie wat gebruik word in die beplanning en
implementeering van 'n binnenshuise konfigurasie-beplannings-
fasiliteit.

'n Oorsig van multiklas-toustaannetwerkmodelle en die SNAP pakkie
word gegee. Die teorie van stogastiese toustaannetwerke word egter
nie volledig verduidelik nie.

Vervolgens word die verskillende fases van die projek behandel.
Probleme wat ondervind is met die versameling van werkverrigtingsdata
word bespreek.

Die probleem van werkladingkarakteriseering en hoe om 'n bevredigende
klassifikasiestrategie vir GFSA te ontwikkel word dan behandel.

Die ontwerp, kalibrasie en geldigheidstoetsing word aan die hand van
die GFSA-model bespreek. Die gebruik van die model vir die doel van
voorspelling word verduidelik met behulp van 'n paar voorbeelde.

Laastens word die geheuebestuurmodel bespreek - hoofgeheue vorm nie deel van die SNAP-model nie en moet as 'n afsonderlike onderwerp hanteer word.

DECLARATION

The subject matter of this dissertation is taken from a joint project
carried out between Gold Fields of South Africa Limited (GFSA), and
the Institute of Applied Computer Science (ITR) at the University of
Stellenbosch. The terms of the project were for GFSA staff to perform
all the necessary work, only calling on assistance from ITR when
required. It is with regard to this project status and due to the
fact that I am the only member of the GFSA staff assigned to the
configuration planning project, that I am able to declare that the
compilation and content of this dissertation is my own work. I
nevertheless fully acknowledge the help and guidance I have received
from my associates, both at GFSA and ITR.

CONTENTS

Chapter

LIST OF ABBREVIATIONS

CC   Communications controller

CR10   Card reader

DFC   Disc file controller

EDS   Disc drive

GFSA   Gold Fields of South Africa Limited

GPC   General peripheral controller

ICL   International Computers Limited

ITR   Institute of Applied Computer Science

JMT   Journal message type

LP   Line printer

MC   Magnetic tape controller

MPL   Multiprogramming level

NRP   New report program

OCP   Order code processor

OP10   Operator's console

RIRO   Roll in roll out

SAC   Store access controller

SMAC   Store multi-access controller

SNAP   Stochastic network analysis package

VM   Virtual machine

VSI   Virtual store interrupt

LIST OF TABLES

LIST OF ILLUSTRATIONS

Page

Figure

1. INTRODUCTION

1.1 COMPUTER CONFIGURATION PLANNING

The level of sophistication of computer technology is rising at an ever increasing rate. The manager of a large computer installation is faced with the unenviable task of trying to determine the extent of the processing power required for his company's needs. More and more he is looking towards some formal methodology or tool to assist him in his decision making.

The efficient management of a large computer system is no easy task. Firstly, the system is made up of so many components that it is difficult for any one person to fully understand all their functions. Secondly, these components are interconnected to form a large number of co-operating subsystems which cannot be managed in isolation, since the failure or inefficiency of any one subsystem can affect the rest of the system.

In order to understand and control such a system, the analyst can make use of the process of modelling. Work in the area of computer performance modelling has advanced rapidly in recent years. It involves identifying the important components of a system and studying the inter-relationships between them. A mathematical model then expresses these relationships in the form of equations which can be solved to produce various performance indicators.

The theoretical development has emerged in the form of a number of commercially available modelling packages. This dissertation explains how such a package was installed as the basis of an in-house computer configuration planning facility at the Head Office of Gold Fields of South Africa Ltd (GFSA), Johannesburg.

1

1.2 THE GFSA-SNAP PROJECT

The need for a configuration planning facility at GFSA was recognized as long ago as 1979. Many avenues were investigated. This included the evaluation of a number of computer modelling packages. Finally, the decision was taken to install one of these packages - the Stochastic Network Analysis Package (SNAP).

SNAP was developed by a university research institute, namely the Institute of Applied Computer Science at the University of Stellenbosch, Stellenbosch, South Africa. It is a multiclass queueing network analyzer, i.e. the jobs that make up the computer workload can be divided into a number of groups or classes depending on the amount of each system resource they have utilized. This makes SNAP a most versatile modelling package.

The main reason behind GFSA's decision to install SNAP, as opposed to any other package, was due to the fact that SNAP was currently being installed at another Johannesburg site. This site employs the same make of computer as does GFSA, namely an ICL 2900 series computer. SNAP had never before been installed on such a computer. GFSA was assured that all the problems associated with an ICL-SNAP installation had been resolved. As it turned out, this was not the case. These problems are discussed further in chapter 4.

In the end, installation was finalized at GFSA prior to the completion of the project at the other ICL site. This resulted in GFSA being the first ICL site to install SNAP as the basis of a configuration planning facility.

This dissertation is intended to demonstrate the methodology employed in installing a package such as SNAP in a commercial environment. After a brief overview of multiclass queueing network models, the various phases of the GFSA project are discussed in detail. Chapter 5 then gives an insight into the

2

way in which the facility may be used for predicting machine performance.

2.   MULTICLASS QUEUEING NETWORK MODELS

2.1   INTRODUCTION

It is not the intention of this dissertation to present the
theory of queueing network models. The emphasis is on the
implementation of a configuration planning facility and not on
the underlying mathematical theory. A brief overview of the
subject is given here in order to clarify what is to follow.
Further detail regarding the various aspects of the modelling
process are given in chapter 4 when discussing the GFSA model.

2.2   QUEUEING NETWORK MODELS

Modelling involves extracting all the important components of a
system and representing them in terms of a mathematical model.
The relationship between the various components must also be
identified and catered for in the model.   Not all of the system
components justify being included in the design.   The modeller
has to decide which ones contribute significantly to the system
resources, and which ones will only serve to make the model
unnecessarily complex.

A queueing network model is simply a set of inter-connected
queues with their associated service centres. Each service
centre represents some component or aspect of the system where
service can be obtained. Various service disciplines can be
defined, which will determine the manner in which the service
centre can give a customer service. These are explained further
in section 3.

The easiest way to define a model is to present it graphically.

4

A queue/service centre combination is presented thus:
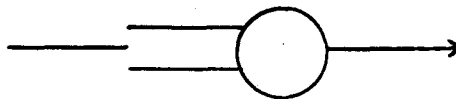
QUEUE   SERVER

The service given by a centre may be such that no queueing ever
results. The service centre is then presented thus:
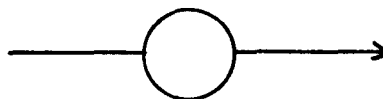
SERVER

Once the network model has been designed it needs to be
analyzed. This is done by expressing the model in terms of a set
of equations and then solving these equations to produce various
performance indicators. Apart from being rather tedious, this
phase is often beyond the capabilities of the average modeller.
However, many computer packages are available commercially which
make model solving a pleasure. One simply has to design the
model, express it in terms of a set of model parameters, feed
these parameters into the package being utilized and wait for
the results of the analysis to be produced.

One such package is the Stochastic Network Analysis Package
(SNAP) developed by the Institute of Applied Computer Science at
the University of Stellenbosch. It employs the stochastic
approach to queueing network model solving. Basically this means
that all arrivals to the various centres are regarded as having
an exponential distribution. The model parameters do not specify
the total time spent at each centre by a job passing through the
system, as does the 'operational analysis' approach. Different
solution methods require different model parameters. Section 4.5

explains the model design process in depth and, together with section 4.6, discusses what parameters are required when using the SNAP package.

Naturally, having analyzed the model, we expect some results that reflect the effects experienced by the queues and service centres. The characteristics of a queue will depend on:

- how frequently customers join the queue;
- how much service each customer requires;
- the service discipline of the service centre.

Various performance indices exist to explain the performance of the queue/service centre combination. The most useful indicators are:

- the time each customer spends in the queue;
- the length of the queue;
- the utilization of the service centre;
- the throughput rate of customers through the centre.

2.3 OPEN AND CLOSED CHAINS

Jobs can circulate through the network using an 'open' or a 'closed' chain. An open chain implies that jobs arrive at the network, entering at one or more of the service centres. They proceed around the network receiving service at one or more of the centres and finally depart once processing is complete. For the network to be stable the arrival rate must obviously equal the departure rate. However, the number of customers in the system at any one time need not be constant.

Closed chains maintain a constant number of jobs circulating through the network. There are neither arrivals nor departures. The multiprogramming level, or concurrency, is determined by one of the model parameters that define the model.

6

2.4 MULTICLASS QUEUEING NETWORK MODELS

A multiclass queueing network model can cater for more than one 'class' of job. In other words, the jobs that make up the workload can be separated into groups depending on the type of resource usage they display. Each group forms a separate chain in the network. These chains are run simultaneously so that the effect of one group on another is correctly reflected in the model.

The advantage of a multiclass network is that different service rates can be defined at each centre for each of the classes. A set of performance indices is produced for each class as well as for the network as a whole.

Each chain can be defined using a number of classes. This means that a job can obtain one amount of service at a centre, leave it and, some time before returning to that centre, change class and thus receive a different amount of service the second time round.

Just how the multiclass queueing network model is defined will depend on the application that one is trying to model. The point that is being made here is that multiclass models show enormous versatility in the variety of applications that can be modelled. The artful designer is almost limitless in the manner of systems he can implement.

3. THE SNAP PACKAGE

3.1 INTRODUCTION

The Stochastic Network Analysis Package (SNAP) was developed by the Institute of Applied Computer Science at the University of Stellenbosch. It is a multiclass queueing network analysis package based on the stochastic theory of queueing networks.

3.2 SNAP OPERATION

An intuitive understanding of the way in which SNAP operates is as follows. Customers are routed around a network that contains a finite number of service centres. The routing is performed by means of a set of branching probabilities, defined for each class in the network. A customer of class x will therefore have a certain probability of being routed from centre A to centre B for all classes x and for all centres A and B defined in the network. The branching probabilities may have a value of zero. This simply means that customer classes can be routed around a subset of the service centres in the model.

Once a customer has joined the queue at a centre, it remains in the waiting line until selected for service according to the service discipline defined at that centre. SNAP permits four types of service discipline. These are:

- TYPE 1 : First-come-first-served (FCFS). All the customer classes that visit the centre must have the same exponential service time distribution.

- TYPE 2 : Processor sharing (PS). The service capacity of

the centre is shared equally among all the
customers present at the centre. If there are n
customers at the centre, then each receives
service at a rate of 1/n of its own service
rate. Each class may have a distinct service
time distribution. The distribution is general,
but must have a rational laplace transform.

- TYPE 3 : Infinite server (IS). Each customer receives
service immediately upon arrival as if it were
the only customer at the centre. The service of
each customer is therefore independent of the
other customers at the centre. Each class may
have a distinct service time distribution. The
distribution is general but must have a rational
laplace transform.

- TYPE 4 : Last-come-first-served pre-empt-resume (LCFSPR).
Each new arrival to the centre pre-empts the
currently executing customer and pushes it on to
a LCFS stack. The newly arrived customer may in
turn be pre-empted by a later arrival. When a
customer's service is completed, it leaves the
centre and the customer on top of the LCFS stack
regains control of the centre to resume its
service. Each class may have a distinct service
distribution. The distribution is general, but
must have a rational laplace transform.

SNAP requires only the centre type and the mean service times
for each class to be input. SNAP itself produces the service
time distribution from these two parameters.

The concept of open and closed chains described in section 2.3
and the concept of multiclass chains discussed in section 2.4
are both catered for by SNAP.

This brief outline has served to introduce SNAP and the various
facilities it offers. Further insight into the operation of the
package will be gained in chapter 4.

4.   IMPLEMENTATION OF THE SNAP PACKAGE AT GFSA

4.1   INTRODUCTION

Implementation of a configuration planning facility using a
package such as SNAP involves a lot more than simply loading the
SNAP programs onto the computer. We require various parameter
values to feed into SNAP. These values include parameters
defining the present system configuration on the one hand, and
parameters defining the present workload on the other. The
former are relatively easy to determine, whilst the workload
parameters are dependent on the type of monitoring available on
the computer. The ICL 2900 series does have a sophisticated
monitoring facility, but it does not supply all of the figures
required by SNAP. Obtaining these figures proved to be the
biggest headache and almost caused the GFSA-SNAP project to be
terminated. Determining just what was and was not being
monitored turned out to be a major part of the project. It is
therefore dealt with in some detail in the following sections.

Of course, monitoring the performance data is not the only work
required. There are many phases to the implementation of the
SNAP package and they can be broken down as follows:

                - monitoring performance data
                - analyzing the performance data
                - designing the model
                - model calibration
                - model validation
                - memory management model

The memory management model is in fact a separate issue, and one
which cannot be satisfactorily resolved using the ICL virtual

11

machine concepts. It is discussed further in chapter 6.

The above phases are required simply to get the package into a usable form for a particular installation. The actual use of the facility for configuration planning or performance analysis is dealt with in chapter 5.

4.2 MONITORING PERFORMANCE DATA

4.2.1 ICL's journals system

Job processing under ICL's VME/B operating system is done by creating a self-contained unit, called a virtual machine (VM), for each job. This logical address space contains all the code and data structures to make the VM appear as an independent machine. In reality the operating system code is not duplicated but the same physical memory resident copy is mapped by the address translation hardware into the logical address space of each VM, thereby making it appear to be part of each VM. Each VM gathers resource statistics whilst it is running. These statistics may be stored in special files, known as journals.

The information is written to the journals using predefined message formats. The data is logically grouped according to the type of statistics being logged. Each logical grouping may contain one or more message formats, and is identified by a unique journal message type (JMT) number. Message type numbers fulfil two important roles:

- they provide a means of classifying messages;
- they make possible the selective routing of certain types or classes of message to relevant journals.

This latter facility is provided by the following mechanism: When a procedure generates a JMT message to be placed in

12

a journal, it does not specify the journal itself, but supplies only the JMT. By means of a special table, called the journal map, the system determines to which journals the message of the indicated JMT should be written. One message may therefore be written to more than one journal, depending on how the journal map has been defined.

The JMT specifies the logical grouping but does not identify one particular message format (eg. a JMT 16 records file transfer statistics. However, there are many different JMT 16 formats, depending on the type of file involved). In order to determine the exact format we make use of what is known as the PETE text number. PETE stands for Print Extended TExt and is used to select the identifying text that should accompany the statistical information when it is written to the journal. The JMT number and PETE text number together identify a particular message format. A program that processes the journal can therefore make use of these two numbers to identify the nature of the data item. (Examples of JMT message formats are given in appendix A).

4.2.2 <u>Performance figures required for input to SNAP</u>

The performance data required as input to the SNAP package depends to a large extent on the way in which the model is designed. The model itself is discussed in detail in section 4.5. However, for the purpose of this section a list of the parameter values required is given below. Explanations, where necessary, are given after the list.

   - session identity and date;
   - session start and end times;
   - identifying data for each job (usernames, jobnames, and profiles);
   - job start and end times;
   - OCP time per job;
   - VSIs per job;

- number of transfers to each disc drive per job;
- average size of transfers to each disc drive per job;
- number of transfers to each tape deck per job;
- average size of transfers to each tape deck per job;
- number of RIROs per job;
- main store occupancy per job;
- number of records spooled per job;
- average concurrency (i.e. MPL) per workload class;
- number of jobs in the session per workload class.

VSIs are virtual store interrupts - i.e. page faults. They access what are known as secondary storage sites. The GFSA system has two secondary storage sites defined on two separate discs. If the VM quota sizes are set too low, many VSIs will be incurred, causing a definite bottle-neck situation.

RIRO stands for 'roll in roll out' and refers to the moving of an entire VM either into main store (roll in) or out to a secondary storage site on disc (roll out). The RIRO site may or may not be the same as the VSI site. GFSA has a separate RIRO site. All terminal jobs are RIROed once every interaction. All background jobs are only RIROed when there is contention for store.

'Records spooled' refers to records being printed on a printer. ICL's printers are run in the same way as batch jobs, the printer jobs being known as spooler jobs. Only the spooler jobs print records and so only those jobs will collect statistics on the number of records spooled. A user job that lists a file is in fact requesting the spooler to list the file.

The data produced by the VME/B monitoring facilities was found to be inadequate in two areas. The first was concerned with the collection of file statistics, while the second was concerned with the main store occupancy of VM's (i.e. jobs) running in the machine.

4.2.2.1 File statistics

VME/B monitoring only provides a total file transfer figure per job. It does not separate this total amount into those transfers that went to each device (i.e. disc drive or tape deck) or into each file accessed. It also only provides the number of block transfers, and not the blocksize. Without the blocksize one cannot determine the actual amount transferred. In order to model the file traffic to each device for each job, one needs the following information:

- job number;
- user name;
- file name;
- device identification;
- number of blocks transferred;
- block size (in bytes).

It is important to have this information for each job so that one may classify the jobs into behavioural groups (i.e. workload classes - see section 4.4).

4.2.2.2 Main store occupancy

Knowledge of the main store occupancy of each job is required so as to model the paging behaviour of jobs. It should be noted that this is a separate issue and does not form part of the model to be analyzed by SNAP. SNAP represents each entity in the system as a centre providing service for customers one at a time. Memory does not behave in this way. It is a resource shared, and competed for, by a time varying number of jobs. It cannot be successfully included in the model, although a centre is present to represent the action of the

15

scheduler. This is discussed further in section 4.5
(model design) and in section 6 (discussion of the
paging model).

VME/B does in fact produce a main store occupancy
figure. However it is calculated on elapsed time rather
than on OCP time, and hence gives a fairly meaningless
figure (VMs are rolled out of main store whilst idle
and spend a lot of their lives on secondary storage).
The OCP integral therefore gives a much more reliable
indication of the time a VM spent in main store. It was
therefore required that a change be made to the way in
which this occupancy figure was calculated, so as to
represent an integral of the occupancy over OCP time.

Therefore, two changes were required to be made to the VME/B
monitoring procedures. It was not until May 1981 that these
changes were accepted as being adequate for the SNAP package.
Reaching that stage was really a 3-way effort: ITR defined the
required changes, ICL wrote the code to effect the changes and
I determined to what extent these changes had been successful.
This was the most critical phase of the whole project - without
the breakdown of the file transfers we could not have
successfully built a model to represent the configuration in
sufficient detail. The project was very nearly aborted on this
point. The following section therefore looks at the changes in
some detail and explains the manner in which I went about
establishing their degree of effectiveness.

4.2.3 Reps and Mod Amends

One is able to 'bend' the VME/B operating system by means of what are known as 'mod amends' (module amendments) and 'reps' (repairs!). These are simply pieces of code that are inserted each day at load time to modify certain areas of the operating system's code. The additional performance data that was required was obtained through the use of these reps and mod amends.

4.2.3.1 File Statistics

Two types of journal messages are produced when a file is accessed. A JMT 62 is produced at file open and a series of JMT 16's are produced at file close. Two reps and a number of mod amends had to be written.

The first rep was designed to log the VOLUME IDENTITY, BLOCKSIZE and SYSTEM REFERENCE NUMBER (SRN) of the file as a JMT 62. The SRN is reused once a file is closed and in itself does not provide sufficient information to match the JMT 62 with the JMT 16's occurring at file close. Additional information, namely the job number, is required. The job number forms part of the standard message header, which is normally inserted into every journal message. However, this message header can only be inserted above a certain level in the operating system's hierarchy (namely ACR level 5). The JMT 62 is logged by a call at an ACR level less than 5, and hence the second rep had to be written to force the standard message header onto the JMT 62.

Apart from these two reps, module amendments had to be written. These mod amends were applied to each of the following record access mechanisms (RAMs) of VME/B:

17

- the serial RAM

- the indexed sequential RAM

- the hashed random RAM

Three changes had to be made to each of these RAMs:

- there is a macro, SET_FILE_OPTIONS, in VME/B's
  SCL (system control language) which allows a
  user to select whether or not he wants file
  information to be produced in his job journal.
  By setting various parameters he can determine
  the extent and nature of the file information
  produced. As we wanted the file information
  (i.e. JMT 16's) to be produced for all files a
  mod amend was required to ensure that the JMT
  16's would always be produced, irrespective of
  any user setting. (Note that obtaining journal
  messages is a two-stage event: the JMT must
  first be PRODUCED and secondly it must be MAPPED
  (i.e. routed) to the journal(s). We were thus
  able to produce JMT 16's and route them to the
  performance journal without incurring a lot of
  garbage in the users' job journals. Those users
  that wish to have the file information printed
  in their job journals can still accomplish this
  by means of the SCL macro SET_FILE_OPTIONS).

- for multi-section files the logging mechanism
  did not automatically switch on the logging when
  file sections were switched. A mod amend was
  therefore required to ensure that this did
  happen.

- a file is uniquely identified by its fully
  qualified filename together with the absolute
  generation number. However, this information is
  not available to the RAM. But, while the file is

18

open, it is uniquely identified by its SRN, and this is available to the RAM. The mod amend therefore inserts the SRN into one of the JMT 16 messages which are produced when the file is closed.

Determining the success of the changes

There were two aspects to determining whether the reps and mod amends were performing their required functions correctly. The first was concerned with whether they were actually logging the data, message headers, and so on to the relevant messages. The second concern was, given that the logging mechanism was working, whether the figures being logged were in fact correct.

The first aspect was easily verified. One simply had to look at the messages being logged to a performance journal to see whether the data was in the right place. This was indeed the case.

Checking the validity of the figures was somewhat more complex. A series of short sessions had to be run, in a very controlled environment. This meant learning how to operate the computer, so as to be able to carry out the tests in the middle of the night when no one was around to interfere. The sessions had to be short so that the messages in the performance journals would be easy to follow. All the jobs that were run, interactively and in background, had to be carefully recorded. The sizes of the files that were transferred to disc or tape, or loaded into main store, had to be predetermined, so as to be able to verify the transfer figures reported in the performance journals. Needless to say, tests had to be performed on all types of files, using all the utilities that could invoke file transfer activity (for

19

example: COPY utilities, EDIT routines, facilities that access the catalogue such as library lists, and so on).

The multitude of tests that were performed will not be presented here. What is of importance are the results and their impact on the monitoring as a whole. The data that was recorded was deemed to be of a sufficiently high standard for the project to continue. However, we need to be aware of the areas where it was not adequate and the extent of the inadequacy.

Observations involving file transfer statistics

The first fact to become apparant was that all loader activity was assigned an SRN of 1 in the JMT 62 at file open. To add to this sad state of affairs, no loader activity ever produced any JMT 16s at file close to say how much had been transferred. The JMT 62 does log the volume identity, and so we are able to proportion the loader accesses amongst the volumes (taking one JMT 62 to represent one access).

Two utilities were found not to produce any JMT 16s at file close. They are the COPYLIBRARYTOTAPE and the DUPLICATEBLOCK utilities. The former is very seldom used, so this slight deficiency could be ignored. However, the DUPLICATEBLOCK facility is used fairly often and is thus more serious. The reason that these two utilities fail to produce JMT 16s is that they do not make use of the RAMs when accessing files. They are part of a series of older utilities that are being phased out. Although this discovery seemed to represent a major setback at first, it was found that none of the other utilities that were tested displayed this type of behaviour. A significant number of DUPLICATEBLOCKs are performed each day. However, when viewed in terms of

the total traffic, it was felt that this inadequacy would be acceptable for the level of accuracy required.

There was one more inaccuracy related to the amount of traffic transferred. When one lists a file or lists the contents of a library to a terminal, one has the option of 'breaking out' or 'quitting' the list before it has completed. When this happens no JMT 16s are produced to record the traffic that was incurred. It is difficult to quantify just how much traffic we could be losing in this manner. It was decided that this inaccuracy could be ignored.

Determining job transfer figures

The state of the monitored data therefore dictated that the transfer figures for each job had to be obtained from three sources. A loader transfer figure, a file transfer figure (for the instances where the reps and module amendments functioned correctly) and an 'other' transfer figure (to cater for those transfers not recorded by the reps) would need to be determined. (Refer to appendix A for a description of the JMTs mentioned below).

- Loader transfers

    The JMT 9 gives a breakdown of the total disc transfers for the job. This breakdown allows one to determine a total loader block transfer figure. This figure is then split amongst the disc drives using the JMT 62 loader access proportioning described above. The blocksize may or may not be taken into account as well. GFSA users tend to use the default 2kb blocksize for most applications, and so only the access figure

21

was used in developing the GFSA model.

- File transfer figure

The file transfer figures are those figures that
are successfully recorded in the JMT 16s at file
close. The JMT 16s are matched with the JMT 62
to determine which volume was involved. The JMT
61 gives the drive on which that volume was
mounted (see appendix A).

- 'Other' transfer figure

The remaining transfers which did not record the
JMT 16s can be obtained from the JMT 9 produced
at job end. As with the loader figure, this is a
total transfer figure for the job. It has to be
split according to the JMT 62 proportioning.

For each workload class it is therefore possible to
determine a total transfer figure for each disc drive.

4.2.3.2 OCP-integral

Figure 4.1 below illustrates the storage occupancy of a program during its execution lifetime in the computer:
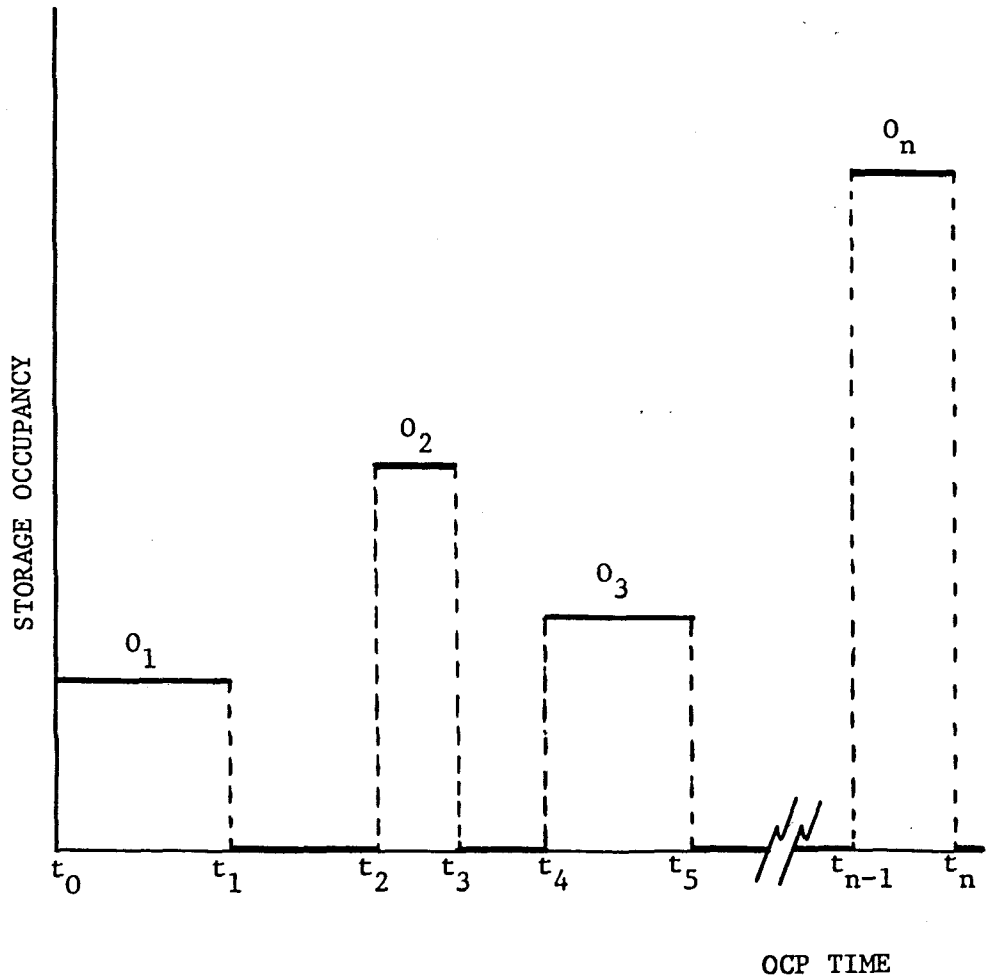
Figure 4.1   Program storage occupancy

The average storage occupancy of the program while in execution is calculated by weighting a particular storage occupancy value by the OCP-time for which that occupancy lasted. It can be seen in figure 4.1 that

occupancy O_1 lasted from t_0 to t_1, O_2 from t_2 to t_3, and so on. In other words the occupancy O_1 lasted for a length of time $t_1 - t_0$, O_2 for a time $t_3 - t_2$, etc.

The OCP-time occupancy integral is thus simply:

$$I_{OCP} = O_1 * (t_1 - t_0) + O_2 * (t_3 - t_2) + \ldots O_n * (t_n - t_{n-1})$$

The total OCP-time is:

$$T_{OCP} = (t_1 - t_0) + (t_3 - t_2) + \ldots (t_n - t_{n-1})$$

The average execution occupancy, which is the quantity required, is computed from

$$\text{Average Occupancy} = I_{OCP} / T_{OCP}$$

This change was effected fairly easily by writing one rep. Problems of overflow and levels of significance were experienced, but eventually a version was produced which seemed to be stable. However, trying to prove that the rep produces a correct figure is virtually impossible, as is discussed below.

Determining the functional correctness

Although each type of VM is given a certain quota, this is only a minimum guaranteed size for that VM. When the machine is busy each job will be restricted to its set quota. If this size is not adequate it simply means that a number of VSIs will be incurred, swopping pages in and out of main memory as they are required. On the

other hand, if the machine is idlc, or there are very few jobs being run, then those VMs that are busy executing are permitted to use up any main store that is available. Pages will only be swopped out when no memory remains unused.

This may make for a very efficient main store strategy. However, it makes it very difficult when trying to determine whether a main store occupancy monitor is functioning correctly or not.

The method employed to test the occupancy rep was as follows. A fortran program was written to initialise a large 2-dimensional array to zero. The compilation listing showed that the program size was 402kb.

The program was first run during the day while the machine was fairly busy. Many runs were executed. The occupancy figures obtained ranged from about 200kb to 500kb. The VM size was set at 216kb. Of course, these figures raised a lot of doubt as to the validity of the measurement.

The test was then run at night in a dedicated machine. Running the program on its own (i.e. no other programs in the machine, apart from the operating system) produced figures as high as 1800kb. It seemed almost unbelievable that so much extra software should be drawn into the VM to run what was a very simple program.

The final test was to run four versions of this program. They were different in that the array sizes were varied. The resulting program sizes were 200kb, 400kb, 600kb and 800kb. When run singularly they again produced what seemed to be meaningless figures. However, when a number of each version of the program

25

were run together (a maximum concurrency of 12 jobs was
used), the figures produced represented the respective
quota size of the profile that the job was running in
(a number of different quota sizes were used). This was
the first really positive evidence that the main store
occupancy figure was working correctly.

The occupancy figure only affects the paging model (see
chapter 6). It does not affect the basic model to be
analyzed by SNAP. In view of this, it was decided to
accept the figure as a reliable indicater until proved
otherwise.

4.3 ANALYZING THE PERFORMANCE DATA

4.3.1 Requirement

Once the performance data has been collected in the performance journal, it needs to be analyzed and the parameters extracted for input to SNAP. Basically this entails writing a program to read the performance journal, extract the relevant information and print it out. There is no great mystery surrounding this program. However, the facilities it provides are important to what follows. Therefore a brief explanation of its facilities will be provided here.

4.3.2 New Report Program

The New Report Program (NRP) was written by ITR. ICL do have a couple of programs that analyze the performance journal, but these were not adequate in that they obviously do not cater for the information provided by the reps and mod amends. Hence the 'new' in New Report Program.

The NRP is only applicable to ICL machines running under VME/B. However, the ideas are universally applicable. The NRP consists of a set of COBOL programs that read the JMT messages in the performance journal and produce the following four reports:

4.3.2.1 Summary Report

This report gives a one line entry for each job run in the session being analyzed. It includes the following data;

- user name
- job name
- profile

- job start time

- elapsed time

- ocp time

- ocp time including system overhead

- machine speed in thousands of instructions
  per second

- vsis per second

- user pages per vsi

- system page transfers

- file transfers

- user disc transfers

- user tape transfers

- magnetic media transfers per second

- director transfers

- occupancy in kilobytes

- number of riros

Not all of these figures are used in calculating the
model parameters, and further explanation of these
figures will only be given where required.

4.3.2.2 Device Usage Report

The device usage report is concerned with the transfers
to the disc drives and tape decks. For each device, it
provides:

- the device address;
- the volume identity (where removable volumes
  have been used, the identity of the volume
  last mounted is recorded);
- the average blocksize (in bytes);
- file accesses and transfers;
- loader accesses;
- other accesses.

The motivation behind the three types of accesses has been explained in section 4.2.3.

The loader and 'other' transfers are determined from the statistics report. These transfers are divided up according to the percentage of the total accesses that each device represents. This is of course, not strictly true. One should take the blocksize into account as well. However, it was felt that sufficient accuracy could be achieved using a pure percentage split, as was proved to be the case once the calibration stage was reached (see section 4.6).

4.3.2.3 Statistics report

The statistics report lists various figures together with the average and variance for each value. We are interested in the following figures:

- ocp time;
- ocp utilization;
- vsis;
- director transfers - total (used for loader transfer figure);
- I/O transfers - (disc+tape)-files (used for 'other' accesses figure);
- records spooled (for the printer workload);
- nember of riros;
- elapsed time (used when calculating the MAC arrival rate).

4.3.2.4 Concurrency distribution

This report gives the time and percentage of time that the multiprogramming level for this workload class was

at each level. It also gives an average concurrency, and this is the figure we are interested in.

Examples of each of these four reports are given in appendix B.

The NRP has a number of selection facilities. In other words one may select (or exclude) certain jobs from a session based on the following criteria:

- user name;
- job name;
- profile;
- start and end times for a selected session;
- store occupancy - min and/or max values;
- ocp time   - "   "   "   "
- vsi rate   - "   "   "   ",
- I/O rate   - "   "   "   "

By employing various combinations of these selection criteria one can divide the jobs in a session into workload groups. In other words, one can group together all jobs which exhibit similar behaviour (eg. OCP bound jobs, I/O bound jobs) or one can extract all jobs belonging to a certain user, say, to see just what tendencies his jobs exhibit. Obviously there are many combinations that one could try, so it is important to be realistic when classifying the total workload into workload groups. The GFSA workload does not lend itself very well to such classification. A fair amount of effort went into the whole question of workload classification and this is discussed in detail in section 4.4 below.

4.4 WORKLOAD CHARACTERIZATION

4.4.1 Overview

The function of workload characterization is to identify the various workload types and to specify their resource demands. A workload type can be defined as a set of jobs with similar resource requirements. Each workload type is then represented in the network model by a unique workload class. The number of customers in each workload class is set to the concurrency or multiprogramming level (MPL) for that class, as determined by the concurrency report of the NRP (see section 4.3.2.4).

The facilities offered by the NRP have been explained, in particular the facility to select jobs according to certain criteria. Having decided on which jobs to select for a particular workload type, the NRP will produce a summary report, device usage report, statistics report and concurrency distribution for the subset of jobs that qualify for that workload type. A series of selections is made, one for each workload type to be represented in the functional model. So, once one has decided on the criteria to be used in selecting the workload types, it is a simple matter to obtain the workload parameters for input to the model. The problem, obviously, is to define just what is to constitute each workload class. This is no easy task. It is influenced by such diverse items as:

- hardware being used;
- software being used;
- jobname standards;
- programmer discipline;
- office hours.

I delved quite deeply into the whole question of workload characterization, using the BMDP statistical package to plot

scattergrams, histograms and cluster analyses. The results were
fairly unsatisfactory in that most users performed a variety
of work (eg editing, compilations, program runs - in background
as well as via a terminal) but usually used the same job code
when logging in to the system, thus giving no indication of the
actual work performed by that job. Below is a summary of my
findings which indicate the unsatisfactory conclusion I was
forced to reach - namely that for the purpose of building a
GFSA model, we could not effectively break the workload down
any further than into three classes. These are a BATCH workload
class, a terminal (MAC) workload class and a SYSTEM workload
class comprising such jobs as the printer, card reader and
operating system control jobs. (MAC is short for MACRESPOND and
stands for multi-access-responder).

4.4.2 GFSA workload characterization findings

4.4.2.1 Bottom-up approach

This approach implies arranging the jobs into subsets
according to their resource usage. One would typically
take a daily session, although a group of sessions may
show more clearly defined clusters. On the other hand,
it may be that the workload varies depending on the day
of the week, so one would need to be careful that one
does not obscure such daily trends when combining many
sessions together.

Basically, the bottom-up approach involves:

- deciding which measurement to use in forming the
  sets (eg. file transfers, OCP time, VSI rate,
  OCP/elapsed time, or any combination of the
  available measurements);

- deciding on the number of sets (classes) to be created;

- producing a list of the allowed ranges for each measurement appearing in each set, as determined by the cluster analysis package, to be used as selection parameters for the NRP.

The bottom-up approach would produce the most accurate classification of different workload types. However, I see it more as a once off exercise to be used for the purpose of performance improvement, rather than in a configuration planning exercise. It assumes that we accept that the resources being utilized are necessary and makes no attempt to identify just which jobs or users are responsible for bad machine usage. One can of course do a separate analysis to see which jobs were responsible for each type of resource usage, but what is implied here is that it should be possible to do the workload characterization dynamically and to prepare the parameters automatically for input to SNAP.

4.4.2.2 Top-down approach

The available data must be analyzed in an attempt to define various sets. However, in this approach, these sets must be labelled by job identity information (and not by the data as in the above approach) such as:

- user name
- user number
- system code
- accounting code
- job code
- job identification

(The GFSA jobname standards are described in appendix C).

33

One could also use some broader form of set labelling
such as a division between background and interactive
terminal jobs. Or a wider split for the background
jobs could be achieved using the batch profiles (see
appendix D). Another possible division could involve
the various sections within the computer division. GFSA
has the following sections: ISD (Information Systems
Development), OR (Operations Research), SW (Software,
or Technical Support). This division could be achieved
using parameters in the job name. Once the set labels
have been defined, it is a relatively simple matter to
separate the jobs into their respective classes using
the NRP.

This approach has an additional advantage. It is fairly
easy to model the effect of a new system on the machine
resources. We will need to know enough about the new
system in order to classify it according to our
predefined sets. It could be inserted into a set in
three ways:

- increase the concurrency of that class in
  the model;
- duplicate a job already in that set;
- try to estimate resource usage (this will
  maybe require discussion with the systems
  analyst) and then insert a dummy job into
  the data.

4.4.2.3 Application of the above theory to the GFSA system

As mentioned above, the bottom-up approach would give
the most accurate results in terms of response time and
showing where the bottlenecks occur. One could build a
model and make predictions concerning changes to the
configuration, but it would not be as easy or realistic

to make changes to the workload parameters. When one is
working in a commercial environment, one has to be able
to relate one's figures to managerial understanding.
This implies making changes to the workload in terms of
jobs or systems being run on the machine, not purely in
terms of the resources being used. VSI's and RIRO's are
Greek terms so far as managers are concerned (I should
explain that managers in GFSA are moved from division
to division in order to gain experience in all aspects
of the company. This does mean, however, that the DP
manager is very often not from a DP background). The
bottom-up approach is certainly a valid approach, but
not in the GFSA context. Nevertheless, I did run an
exercise to extract all the high resource usage jobs.
The resources used were: OCP time, virtual store disc
transfers, magnetic tape transfers, VSIs and main store
occupancy. The following results were apparant:

- the line printer and master operator jobs were
  present in most runs. This is deceptive as these
  jobs run for the whole session and therefore
  their readings are bound to be high.

- table 4.1 shows those usernames and/or jobnames
  that could be associated with high resource
  usage. The figures given indicate the percentage
  of jobs labelled with that username/jobname
  combination that displayed the resource usage
  indicated at the top of each column, eg. 16% of
  all CCGEOLDB jobs used more than 150 seconds OCP
  time.

  Some job types could be specifically identified
  eg. :SYSTEM.B8140PRPS4070. This type of job
  would be easy to separate into a specified work-
  load class.

35

JOB IDENTIFIER	TOTAL NO OF JOBS	OCP TIME GREATER THAN:			VSIs GREATER THAN:			TAPE TRANSFERS GREATER THAN:			VIRTUAL STORE DISC TRANSFERS GREATER THAN:			MAIN STORE OCCUPANCY GREATER THAN:		
		150	250	500	2500	4000	9000	5000	6000	10000	30000	50000	75000	50	100	150
CCGÉOLDB	473	16,0	11,1	7,4	13,1	5,1	0,0	1,5	1,5	1,3	20,0	10,9	6,1	7,4	1,5	1,1
CSSTFR	404	12,1	7,7	3,7	10,9	4,2	0,0	11,1	9,7	3,5	18,9	11,0	6,9	3,7	0,5	0,0
ENGINEER	227	4,2	0,9	0,4	1,3	0,4	0,0	12,2	11,5	6,6	15,1	11,8	6,2	0,0	0,0	0,0
ORANB	92	10,9	8,2	4,3	2,2	0,0	0,0	0,0	0,0	0,0	13,8	9,1	6,8	1,1	0,0	0,0
ORORE	244	14,3	11,3	4,9	5,7	0,8	0,4	0,0	0,0	0,0	34,2	22,7	16,3	1,6	0,4	0,0
ORROCK	96	12,5	11,5	9,4	3,8	0,0	0,0	0,0	0,0	0,0	19,7	11,2	10,7	2,1	0,0	0,0
ORSTATS	98	5,4	3,4	1,1	13,3	5,2	2,1	4,3	3,0	3,0	11,3	7,6	3,4	3,2	2,1	0,0
MISBP	315	35,5	27,5	16,5	20,7	13,3	0,0	41,4	29,1	25,3	2,6	1,3	0,6	23,9	12,9	1,3
MISBP.....UP	106	76,6	58,6	43,2	52,4	33,9	0,0	73,0	69,2	59,4	0,0	0,0	0,0	62,8	30,2	2,8
SYSTEM.B8140PRPS4070	43	88,7	88,7	86,3	42,6	30,8	0,0	58,5	56,9	56,9	0,0	0,0	0,0	91,2	91,2	91,2
SYSTEM.B1000P	311	0,0	0,0	0,0	0,0	0,0	0,0	83,6	80,8	62,5	0,0	0,0	0,0	0,0	0,0	0,0
TOTAL JOBS IN SAMPLE	13273	1119	626	251	1352	672	163	627	561	416	2225	1311	762	589	208	81
JOBS IN SAMPLE / TOTAL JOBS (%)		8,4	4,7	1,9	10,2	5,1	1,2	4,7	4,2	3,1	16,8	9,9	5,7	4,4	1,5	0,6

TABLE 4.1   Resource usage for the period 6/11/80 - 10/2/81

Figures shown are percentages : No of jobs in classification / Total jobs of that identity

However, in most cases the username was the only criteria that could be used for identification, eg. there was nothing in the job identifiers of the 16% of CCGEOLDB jobs mentioned above that could be used to differentiate these jobs from the other 84% of CCGEOLDB jobs. In other words, if we decide that the percentage of CCGEOLDB jobs showing high OCP usage is great enough to warrant a separate class of CCGEOLDB jobs, then we would have to include all the CCGEOLDB jobs in the class. (If we checked the actual OCP time for each job before assigning it to a class, we would be reverting to the bottom-up approach).

There are three aspects to consider when deciding on a separate class:

- the cut-off point for the data (150 seconds in this case);

- the percentage of jobs of this type greater than the cut-off point (16% in this case);

- the size of the sample (8,4% in this case).

It should also be noted that table 4.1 only shows nine usernames. GFSA, at the time when this analysis was performed, had defined 82 usernames. In other words only 11% of the users showed some form of abnormal resource usage. The other 89% all showed similar job tendencies. Obviously these percentages are directly related to the sampling levels taken in the analysis. Nevertheless, it was clear that the bottom-up approach was unsuited to the type of class formation we were aiming at.

The top-down approach has the advantage of being

37

intuitively sound so far as management is concerned. There are a few clear cut divisions that one can make, and it is a simple matter to select the jobs falling into each division using the selection facilities offered by the NRP. For example, the two character system code (see appendix C) can be clearly related to the various sections within the computer division. We could therefore select jobs according to the sections by selecting on system code. Whether this is a valid selection to make is of course a totally different question.

The GFSA system allows for the following top-down divisions to be made:

- Profiles. The five background profiles (see appendix D) would form an ideal basis for dividing the batch work. I use the word 'would' advisedly. The problem is that the profiles are rarely used in the manner for which they were designed. The programmers tend to use whichever profile is available. A further point is that we are only interested in the batch workload in so far as it affects the terminal response times. This is really in the area of performance tuning which is done by adjusting the scheduling parameters. It would only serve to make the model unnecessarily complex if one were to set aside five workload classes purely for the batch workload.

- Sections. As I mentioned earlier, this is a nice intuitive division to make. However, in practice I found that there were not many differences in the workload produced by each section. The exercise performed here was to divide the data into the following sections: SW (software i.e.

technical support section), ISD (information systems development), OR (operations research), production jobs and 'other' (eg. interactive programs run by the Financial Division). The jobs in each of these five sections were then further subdivided into MAC and batch.

The only differences of any note were:

- high magnetic tape transfers for production jobs and SW batch jobs;
- high disc transfers (i.e. to user files) for OR MAC jobs;
- high magnetic media transfers (i.e. total disc traffic) for ISD batch jobs.

The data used in this exercise spanned a period of five weeks. It can therefore be accepted that the workload is reasonably stable from one section to another.

- Batch, MAC and 'other'. This seems to be the neat way of dividing the workload into classes. I feel that this is still a rather broad form of classification. However, one can use the NRP to select a set of jobs and then investigate such things as inter-VSI time. One does not always have to take the analysis through SNAP. Therefore, for the broader requirements of configuration planning it is sufficient to build the model on these three classes. There is really no point in struggling to form several classes from a workload that does not display varying characteristics.

Apart from the brief reference to production jobs above

nothing has been said about separating production and development work. This is because the GFSA production jobs are run almost totally after hours, and we are not too concerned with this period of the day. The only requirement is that the production runs be completed by 0800 the following morning. Therefore the analysis of the workload is restricted to the period 0800 - 1630. Any production jobs that do 'creep' in before 1630 are treated in the same way as the normal batch jobs.

To conclude, the question of workload characterization has been dealt with fairly thoroughly. However, the method of operation at GFSA does not lend itself to an ideal classification strategy. However, the three work types finally chosen proved to be a satisfactory basis for the functional model, and as discussed in the next section, resulted in rather encouraging calibration results.

4.5 MODEL DESIGN

There are two main aspects to designing a model to represent a computer configuration. These are the physical model design and the functional model design. The former is concerned with representing the hardware configuration of the computer installation as a set of service centres connected in accordance with the physical connections of the real system. The latter involves the functioning of the system, how the components interact, their relationships and any software or workload idiosyncracies.

4.5.1 Physical model design

The procedures used in physical model design are relatively straightforward. One starts with a diagram that includes all the hardware components of the computer system. These are then translated into a queueing network model with a number of centres representing each important component of the system. It is not necessary that every component or device be included in the queueing network model. Some devices experience no processing delay, whilst others, such as control units, do not form part of any of the logical routings in the system. An item such as main memory can be used by a number of jobs at the same time, and in itself does not represent a delay. Jobs may have to wait for store to become available, but this is then modelled by a service centre which represents the scheduler. Once a job is in store and executing, memory plays no part so far as the queueing network model is concerned.

Note that a service centre need not represent a hardware component of the system. The scheduler mentioned above is an obvious example. This means that we may find ourselves adding more centres during the functional design phase (see section 4.5.2).

41

4.5.1.1 GFSA physical model design

The GFSA configuration is shown in figure 4.2. It shows all the major components of the system. Figure 4.3 shows the corresponding queueing network model that represents this configuration. Many of the items in figure 4.2 are not included in figure 4.3 and brief explanations are given below.

The SMACs (Store Multi-Access Controllers) and the SAC (Store Access Controller) do not form part of any logical routing as discussed earlier. They are thus not included in the model.

The timesliced nature of the OCP is represented by a service centre with a service discipline of PS (see section 3.2).

The DFCs (disc file controllers) and disc drives are represented by FCFS service centres. At first glance it may seem odd that the DFCs appear after the disc drives. This is because the delay at the disc drive centre is representing the seek time. The delay at the DFC represents the duration that the channel is busy. In other words, a seek may be performed on one disc while the channel is busy transferring data to or from another disc, and this is the effect we wish to model.

The GPC (general peripheral controller) causes no significant delay and is not included in the model.

There are six tape decks, all connected to two magnetic tape controllers (MC). The system is designed such that a job wishing to access a tape simply chooses whichever MC is available, or becomes available first. The only limiting factor is then the number of MCs and this is modelled by having two service centres each handling
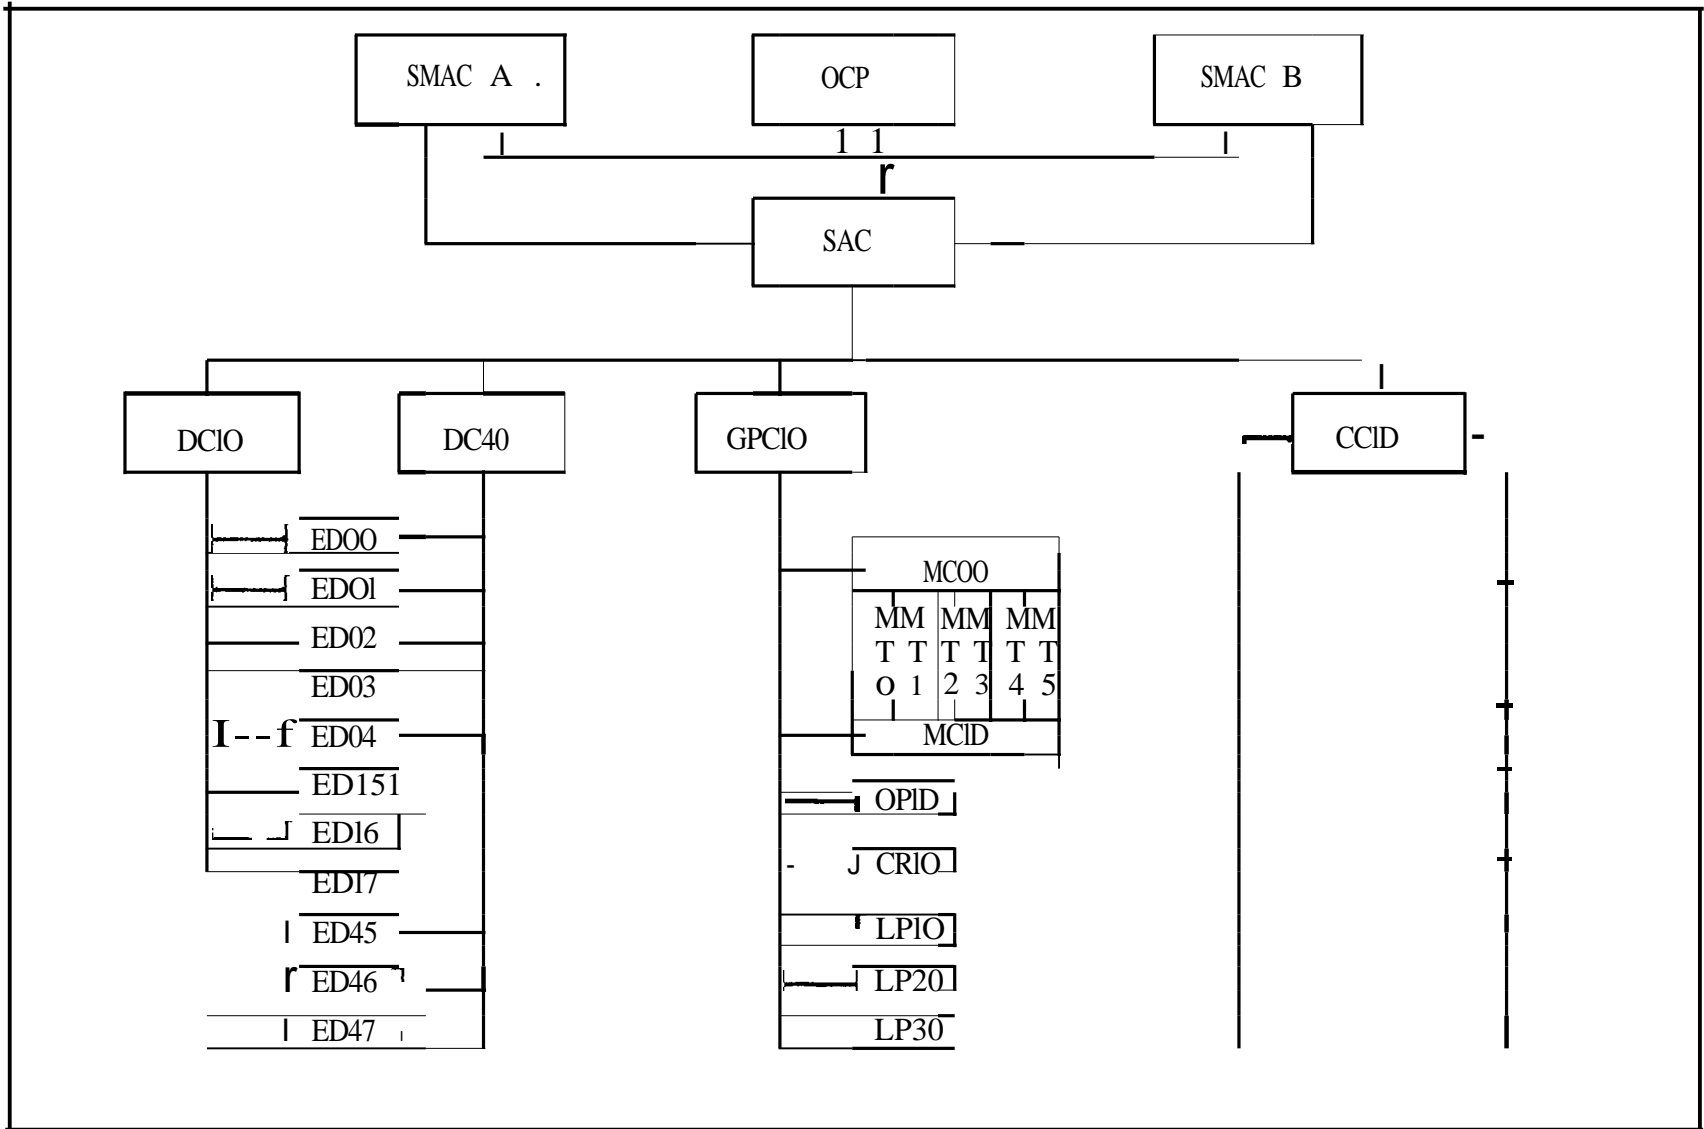
Figure 4.2   GFSA system configuration
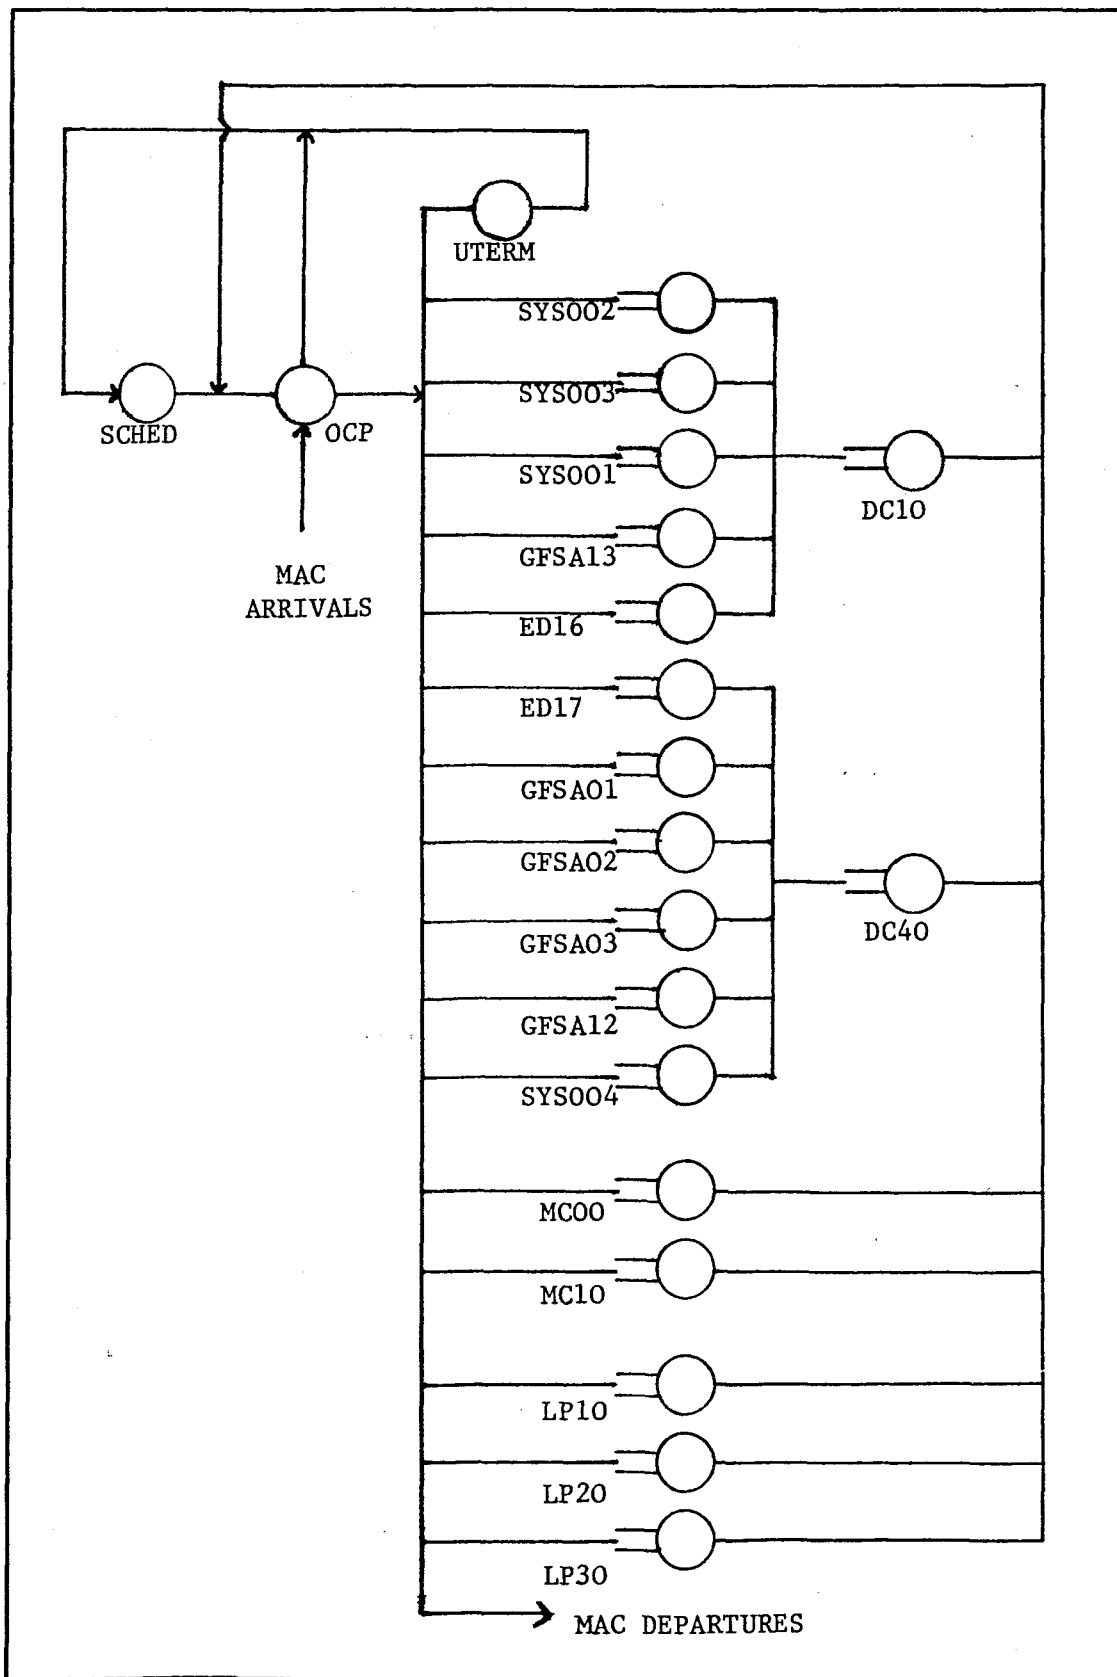
Figure 4.3   GFSA preliminary queueing network model

half of the magnetic tape transfers using a FCFS service discipline.

The operators' station is not included in the model as a service centre. However, the workload imposed by the operators' station is included in that the console is run by a job, called the 'master operator responder' job. This job behaves in essentially the same way as any other job so far as gathering statistics is concerned. It is included in the SYSTEM workload.

The line printers are run by jobs as well. They are fired off from the operators' station in the same way as a batch job would be, and collect statistics in the same way as well. These jobs also form part of the SYSTEM workload.

The card reader is ignored. The statistics are gathered in the normal way, but no provision is made for the card reader job in the model. Its impact on the machine is so slight that it would only serve to make the model unnecessarily complex.

The communications controller and terminals are not modelled explicitly as service centres. The MACRESPOND workload is modelled as an open chain and this is discussed further in the following section.

4.5.2 <u>Functional model design</u>

There are two aspects to functional model design. The first is concerned with the logical routings within the queueing network model, which is really a software refinement of the 'hardware' model discussed above. The second concerns the definition of the various classes to be used. The three workload classes (BATCH, MAC and SYSTEM) each form one chain in the model. Each

chain may consist of one or more classes, and it is the definition of these classes that is of concern here. These two aspects will be explained by describing each of them in terms of the GFSA model.

4.5.2.1 Logical routings

The first addition to make to the model design as presented in section 4.5.1, is to include a termination centre. It will have a routing equivalent to the number of jobs that completed, and will be given a very large service rate so that no delay will be experienced by jobs passing through it. These jobs are routed straight back into the network and serve to maintain the MPL or concurrency at a constant level.

The second addition to be made is to include a service centre to represent the scheduler, or more specifically to represent the delay for jobs waiting to get into main store. This time is not measured. The service rate for this centre is in fact the parameter that we vary during the calibration stage, in order to get the model to produce the same performance indices as given by the measured results. The validation stage will then show whether or not the value set for this parameter is valid.

The model is now essentially complete, and is in fact in a usable form (apart from the class definitions which are discussed below). However, we can make it a lot more 'user friendly' by making some more additions. As it stands at the moment we will need to calculate branching probabilities from the OCP to each disc drive, each magnetic tape controller, each line printer and to the termination centre. If we then wished to

change just one value, for example the number of jobs completing, we would need to recalculate each of these branching probabilities. This is made simpler and more usable by the use of 'micro centres'. These centres are characterized by a very fast service time and so do not invoke any delay. The following micro centres are defined in the GFSA model:

- UTERM : this is the termination centre described above;

- UVSI ; this micro centre represents all the VSIs (i.e. page faults). VSIs access what are known as secondary storage sites. There are two of these sites, on two different discs, so the UVSI centre effectively serves two disc drives;

- UDISC : all the disc traffic, excepting VSIs and RIROs, are routed through this centre. From this centre the traffic can then branch to any of the eleven disc drives;

- URIRO : the RIRO traffic passes through this centre. There is only one RIRO site, so this traffic is then routed directly to the disc drive containing the RIRO site;

- UTAPE ; the magnetic tape traffic is routed through this centre to the tape controllers, with the traffic being equally split between the two controllers;

- UPAPER: the records that are spooled proceed through this micro centre. They are then equally distributed amongst the three printer centres.

The model now contains all the service centres and

routings that are necessary in order to define the GFSA system configuration. The final version of the queueing network model is shown in figure 4.4.

4.5.2.2 Definition of workload type models

As explained earlier, the three workload types that are implemented in the GFSA model are the BATCH, MAC and SYSTEM workloads. In some instances different service rates are required for each workload type. It goes without saying that a separate set of branching probabilities is required for each workload type. The methods used to calculate the branching probabilities and service rates are essentially the same for each workload type, and are discussed in sections 4.5.3 and 4.5.4.

Each workload type is represented by one or more classes in the model. The final model shown in figure 4.4 needs to be refined for each workload type, in some cases even removing certain service centres and routings.

MAC workload

The MAC workload type is modelled using classes 1 and 2. It is modelled as an open chain. An arrival - processing - departure sequence represents one terminal interaction. Processing is as follows: the user presses the SEND button. A short burst of processing power at the OCP is required to initiate 'roll in', i.e. to move the VM into main store. After this a mixture of OCP -processing power, disc transfers and tape transfers is required until the interaction is complete. The VM is then 'rolled out' onto secondary storage and the OCP

48
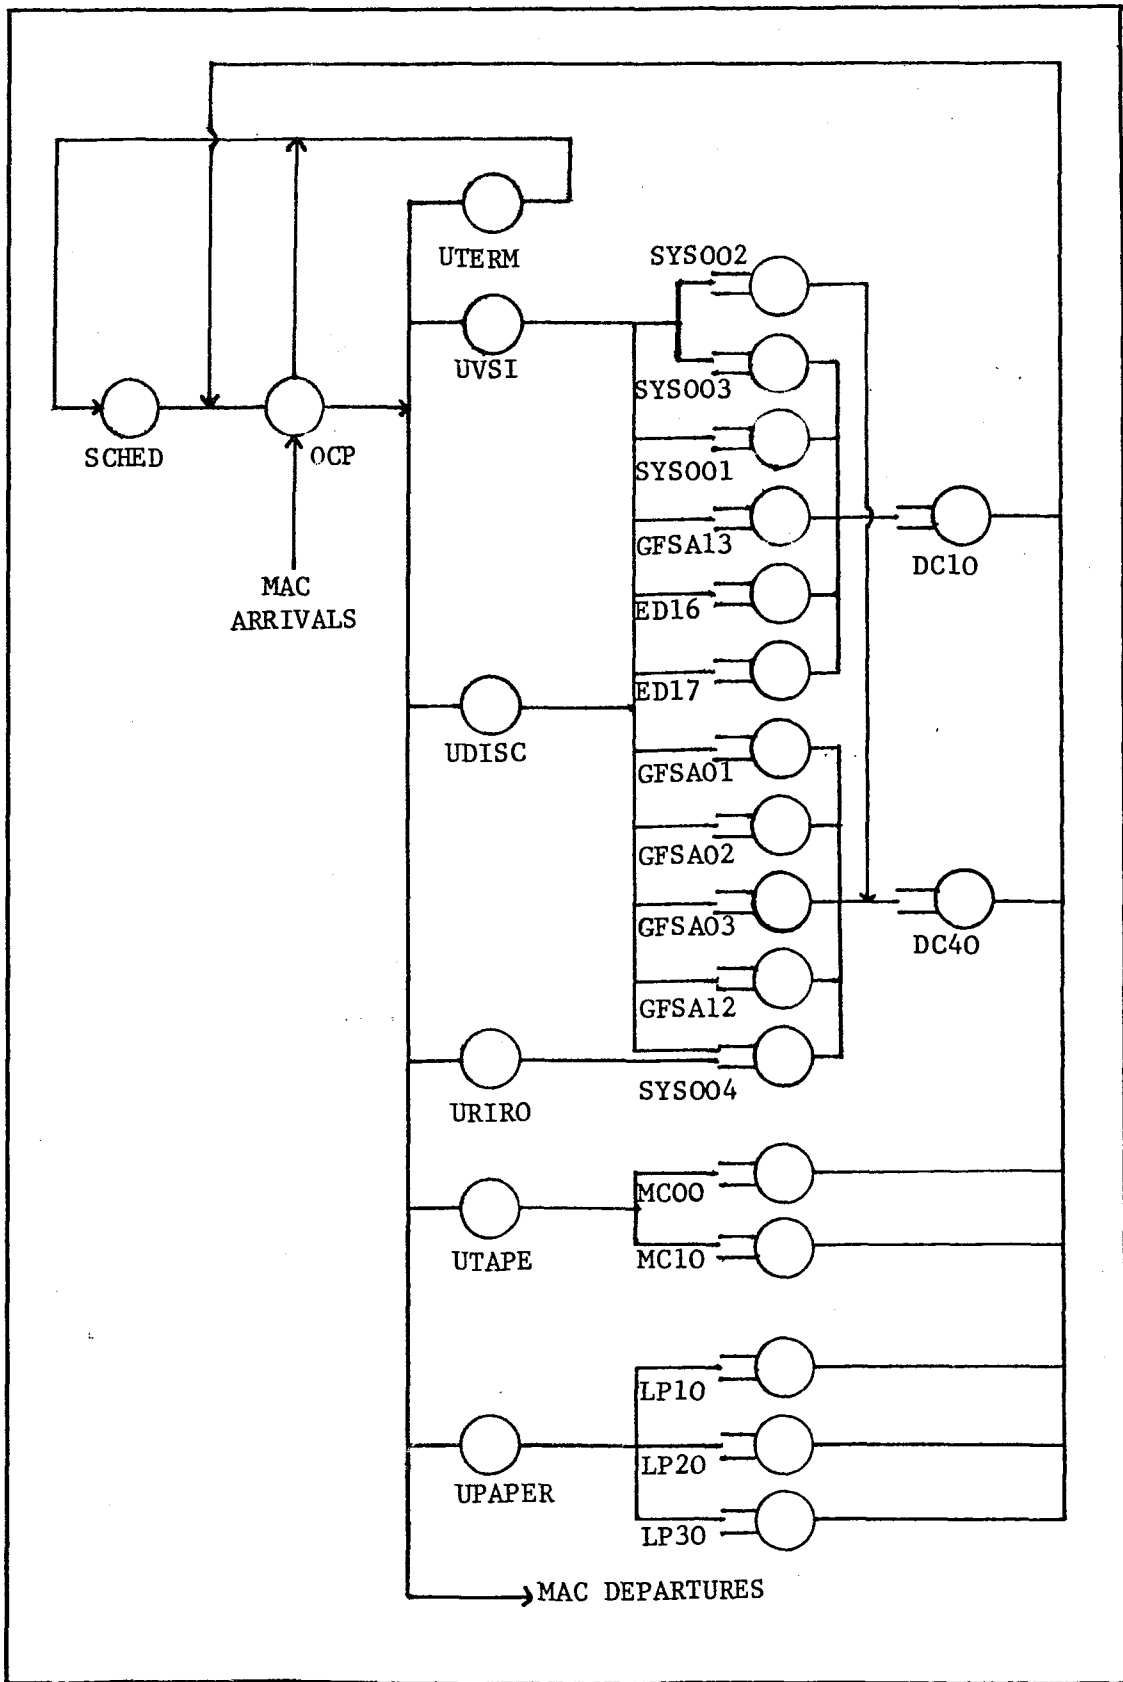
Figure 4.4   GFSA final queueing network model

has a final burst before sending the invitation to type back to the user. (This sequence only includes those aspects that are included in the model. A true sequence would also reflect the action of the communications controller, scheduling priorities, etc. The effect of these items is effectively modelled by the adjustment of the scheduler service rate during the calibration stage).

This processing sequence needs to be reflected in the model. It does not need to be an exact duplicate of the real situation, but must include all the delays that a job can experience. The MAC workload makes use of classes 1 and 2. Processing is as follows: the job arrives at the OCP (calculation of arrival rate is explained in section 4.5.4) as a class 1 job. It proceeds to the scheduler, then to the URIRO centre, on to the disc drive containing the RIRO site and finally to the DFC, all as a class 1 job. At the DFC it obtains double service to represent the roll in and roll out time. From the DFC it returns to the OCP, at the same time changing to a class 2 job. It now proceeds to loop round the network, through the UTERM, UVSI, UDISC or UTAPE subsystems, returning to the OCP each time, until it departs along the DEPART branch. (Obviously the network is solved by means of a set of equations, and this idea of jobs proceeding around the network model is a purely intuitive one. It is nevertheless a valid one). Note that the MAC workload never visits the printer centres. Only the printer jobs access the printers and so only the SYSTEM workload contains these centres.

When there is contention for store we may get what are known as forced RIROs, i.e. the VM is forced out of store before it has finished processing and immediately joins the queue for jobs waiting to enter store. These

50

forced RIROs are included in the total RIRO figure and so are already represented in the model. The number of forced RIROs is a function of the main store size and VM sizes. The GFSA system is scheduled so that very few forced RIROs are incurred.

The final model for the MAC workload is shown in figure 4.5. The arrows and numerals indicate class changes.

BATCH workload

The BATCH workload makes use of classes 3 and 4. It is very similar to the MAC workload in that it uses class 4 for the RIRO routings. It is, however, a closed chain, and so the sequence is slightly different: from the OCP a job may proceed to the UTERM, UVSI, UDISC, URIRO or UTAPE centres. There is a certain probability that it will go to each one (see section 4.5.4). It always leaves the OCP as a class 3 job. If it proceeds to the URIRO centre, the class changes to class 4. It remains a class 4 as it proceeds to the RIRO disc and on to the DFC. It naturally receives a different service rate at the DFC as a class 4 job than it would as a class 3 job. It is delayed for both the roll in time and the roll out time. From the DFC it proceeds to the scheduler, changing to a class 3 job as it does so.

The BATCH VMs only RIRO when there is contention for store. Any contention will first force the MAC VMs to RIRO, so BATCH VMs do not RIRO very often. However, it does not happen so infrequently, so we must make provision for it.

The final model for the BATCH workload is shown in figure 4.6.
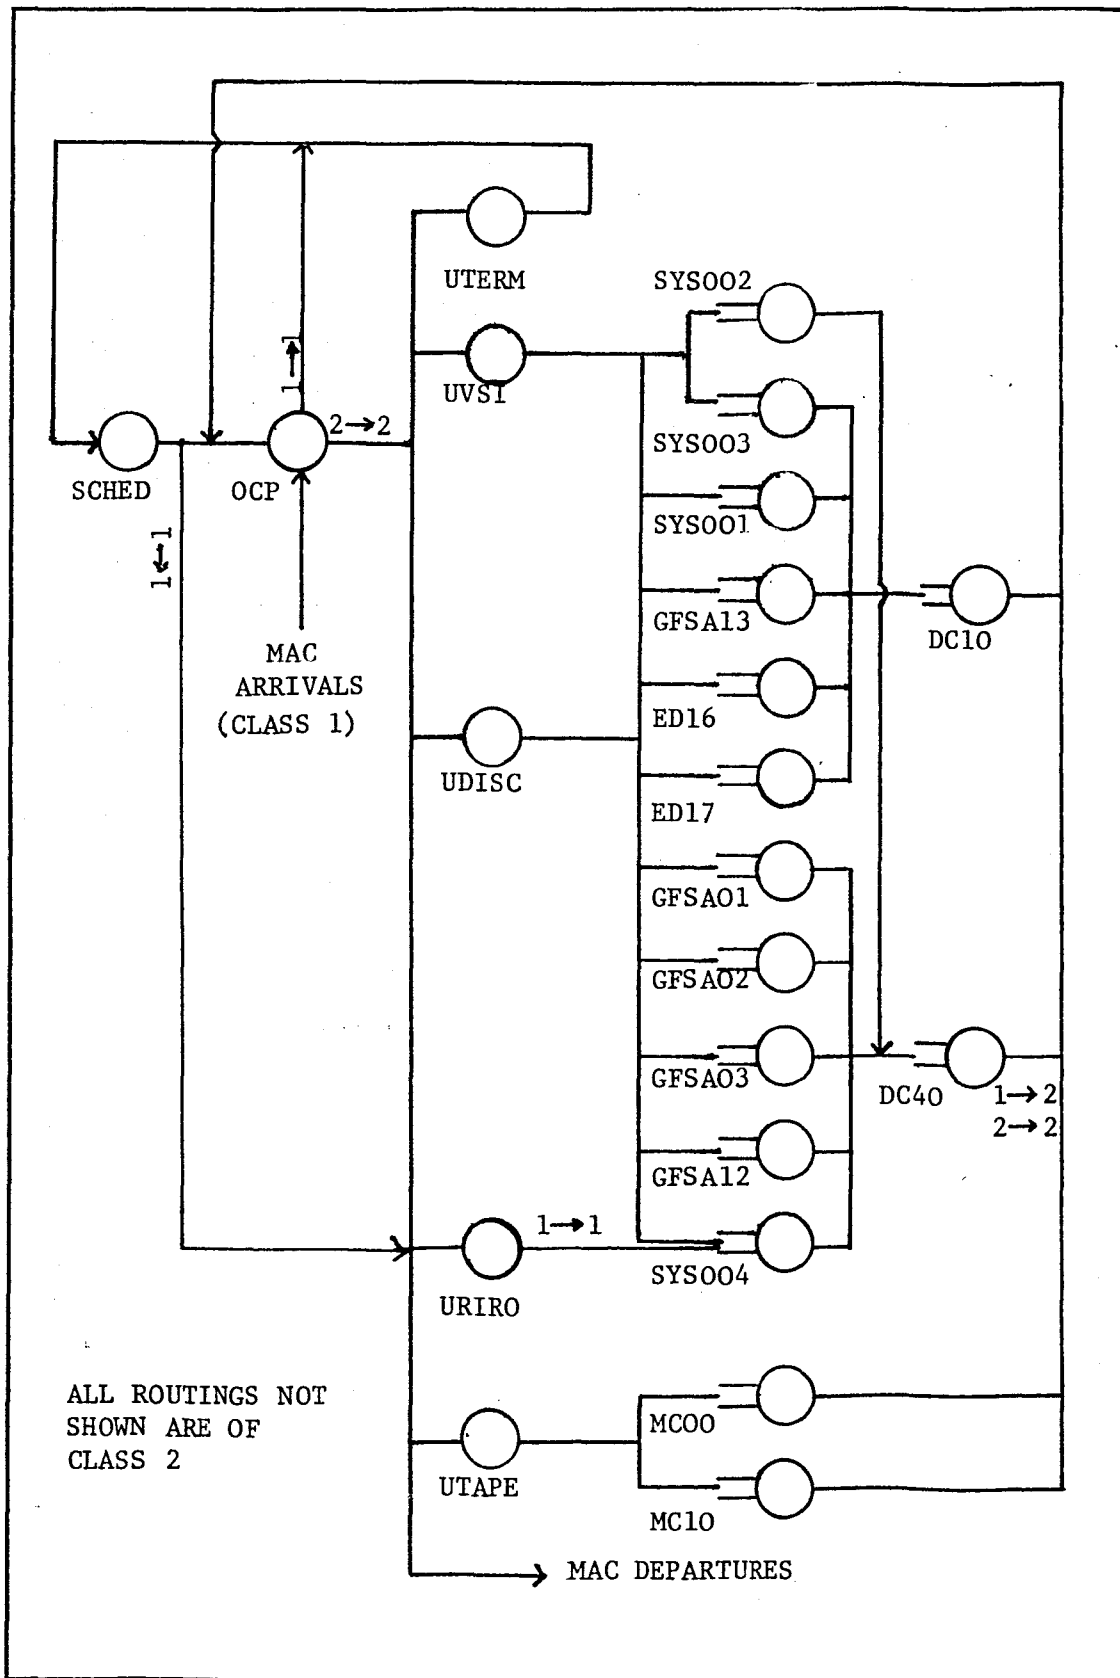
Figure 4.5   Queueing network model for MACRESPOND workload class

ALL ROUTINGS NOT
SHOWN ARE OF
CLASS 3
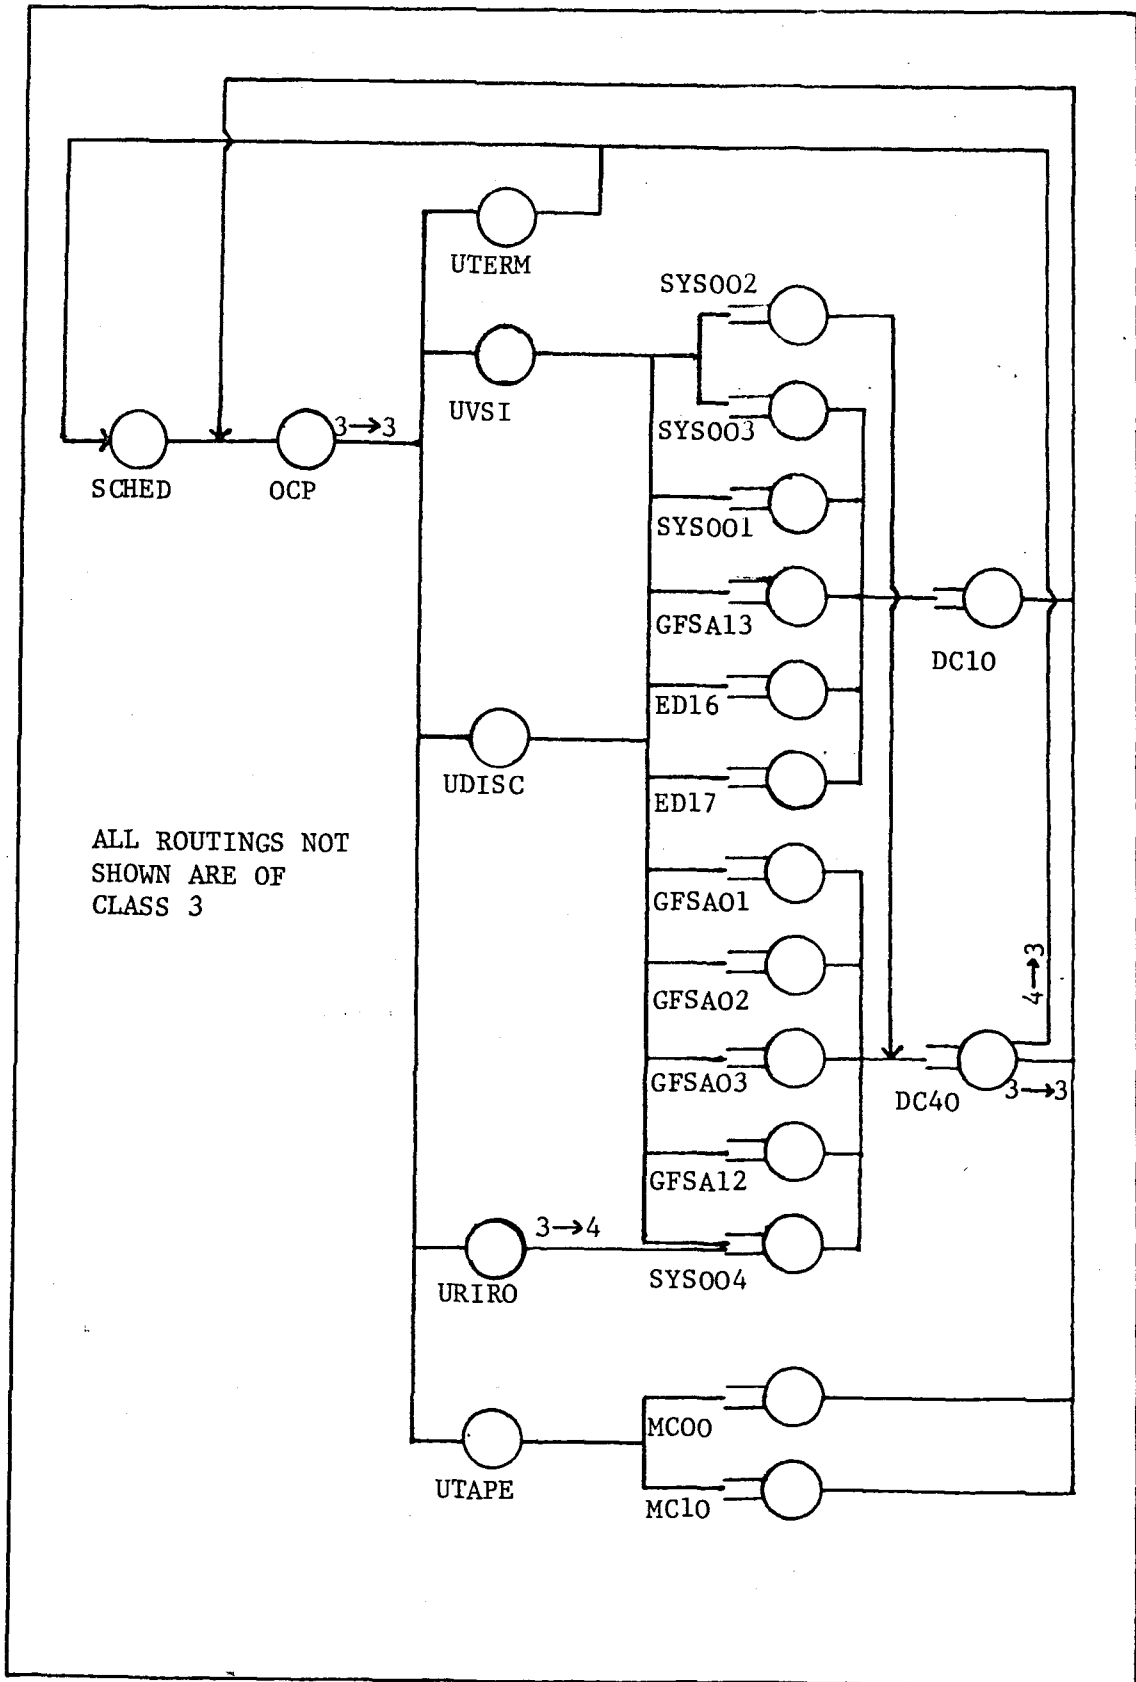
Figure 4.6   Queueing network model for BATCH workload class

SYSTEM workload

The SYSTEM jobs can also be RIROed. However, MAC and
BATCH jobs are RIROed first, so it very rarely happens
that a SYSTEM job will be RIROed. There is therefore no
need to include this possibility in the model. We can
thus model the SYSTEM workload using only one class.
This is class 5.

The SYSTEM model is shown in figure 4.7.

4.5.3 Calculation of model parameters

The model parameters are those values that are fixed, or would
only vary if there were some change in the configuration. They
include such things as number of centres, service rate at each
centre, centre names, number of classes and so on. The model
parameters for the GFSA configuration are as follows:

- number of centres : there are a total of 26 service
  centres in the model (the fact that some workload chains
  may not include all the centres is immaterial here);

- centre names : most hardware units have a specific ICL
  nomenclature. As far as possible, these naming standards
  have been followed.

The disc drives are named in one of two ways. Where a
disc is permanently mounted on a drive, the disc name is
given. Where a drive is used for removable disc packs,
the disc drive identification is given. The permanent
discs are:

          - system discs   : SYSOO1
                                SYSOO2
                                SYSOO3

54
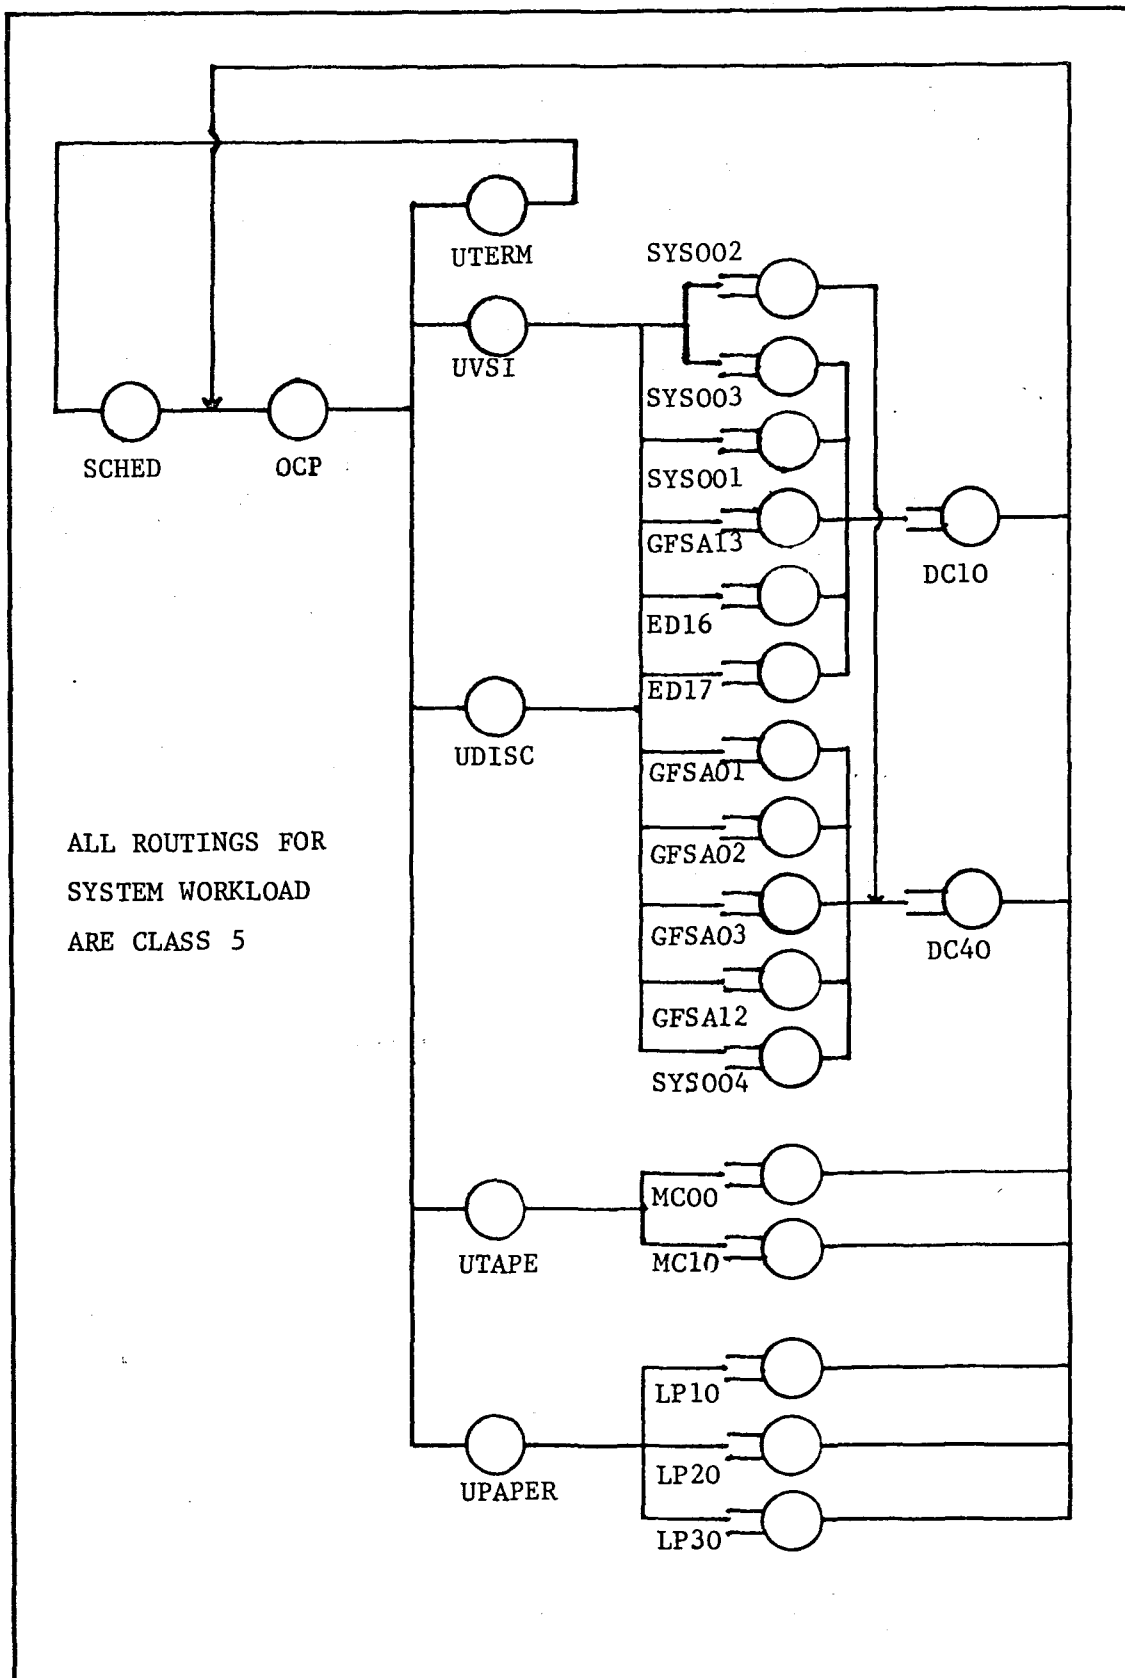
Figure 4.7   Queueing network model for SYSTEM workload class

55

```
                              SYS004
        - user filestore  : GFSA01
                            GFSA02
                            GFSA03
                            GFSA12
                            GFSA13
```

The reusable disc drives are:

```
                            ED16
                            ED17
```

- number of classes : 5 (see section 4.5.2.2);

- service disciplines : the centre service disciplines are
  discussed in section 3.   Only types 1, 2 and 3 are used
  in the GFSA model.

  The following centres have a type 1 service discipline:

      - all disc drive centres;
      - magnetic tape controller centres;
      - line printer centres.

  The following centres have a type 2 service discipline:

      - OCP centre;
      - DFC centres.

  The following centres have a type 3 service discipline:

      - scheduler centre;
      - micro centres.

- service rates : some service rates will vary as the
  workload changes.   These are discussed in section 4.5.4.
  The service rates that are fixed are determined in the

following ways;

- disc drive centres : ICL engineering manuals give a seek time of 30 ms, and a latency time of 8.33 ms - i.e. a total service time of 38.33 ms. The inverse of this gives us the service rate, namely 0.026 per second.

- DFCs ; the disc file controllers are rated at 806 kb/secs. The blocksize can vary, if the users specifically place their files with a different blocksize. The GFSA users tend to be lazy in this regard. Hence a standard blocksize of 2kb can be assumed, giving a service rate of 403 per second.

The RIRO transfers obviously require a different service rate. The MAC VMs have a quota size of 216kb. The 'roll in roll out' delay is incurred during one pass through the DFC (as described in section 4.5.2.2). Hence the model must cater for a transfer of 432kb. With a speed of 806 kb/sec this gives a service rate of 806/432 = 1.8657 per second. The BATCH quota is 288kb for all the profiles except the BMILL profile, which has a quota of 432kb. The BMILL profile is not used very often, so a figure of 288 is used here. This is a total transfer of 576kb, i.e. a service rate of 806/576 = 1.3993 per second.

- MCs : the magnetic tape controllers are rated at 200kb/sec, giving a service rate of 100 2kb blocks per second. Again, the users may define blocksizes other than 2kb, but this is very rare. However, this figure of 100 is not strictly correct. The GFSA daytime work is not very 'tape bound'. It is likely that only one controller is in use for most of the day, giving an effective

service rate of 100. The model on the other hand
will give an effective service rate of 2 x 100.
An arbitrary figure of 2/3 of the service rate
was therefore chosen i.e. each controller was
given a service rate of 67 blocks per second. The
calibration would show just how reliable this
figure was.

- LPs ; the line printers present a slight problem.
  The information we have is as follows:

    - the line printers are rated at 1600 lmp;
    - the lines per minute speed will vary as
      the line length varies;
    - the monitored data gives a figure for the
      number of records spooled;
    - records can be any length;
    - the spoolers have buffers which means that
      there is no hidden delay.

It is therefore fairly difficult to even estimate
an average service rate in terms of records
spooled per second. It was clear that the SYSTEM
workload would need to be calibrated on the
spooler service rate (section 4.6.1 explains the
calibration process). A starting point for the
calibration had to be decided on. By assuming an
average records per minute rate of 1200, the
service rate for the spoolers is then set to
1200/60 = 20 records per second.

- micro servers ; these centres are given a very
  large service rate so as to cause no delay. The
  service rate is set to 9999 per second.

- scheduler ; the scheduler service rate is the
  parameter that is varied during the calibration

58

stage to allign the results obtained from the model with those obtained from the measured performance indices (see section 4.6).

- branching probabilities : many of the routings are considered as workload parameters and are discussed in the next section. Those that are fixed are presented here:

    - branching probability from UTERM : all jobs proceeding through the UTERM centre are sent to the OCP. Therefore, for all classes, a branching probability of 1.0 is given for jobs going from centre 21 to centre 1;

    - branching probabilities from UVSI : traffic to the two secondary storage sites is split equally. There is therefore a branching probability of 0.5 from the UVSI centre to each of the disc drives centres containing a secondary storage site (i.e. SYS002 and SYS003);

    - branching probabilities from URIRO : there is only one RIRO site and it is situated on SYS004. There is thus a branching probability of 1.0 from the URIRO centre;

    - branching probabilities from UTAPE : the tape traffic is split equally amongst the two tape controllers, i.e. a branching probability of 0.5 from the UTAPE centre to each of the magnetic tape controller centres;

    - branching probabilities from UPAPER : the printer traffic is divided equally amongst the three printer centres, i.e. branching probabilities of 0.3333, 0.3333 and 0.3334 (=1.0!) from centre 26

to centres 18, 19 and 20 respectively;

- the disc drives are connected to one DFC as shown
  in figure 4.4. There is a routing probability of
  1.0 from each drive to its respective DFC;

- the DFCs, MCs and LPs are all routed to the OCP
  with probability 1.0, except for the BATCH RIRO
  routing. Class 4 jobs proceed from centre 4 to
  centre 2 (the scheduler) with probability 1.0,
  changing to a class 3 job en route;

- the scheduler has a branching probability of 1.0
  to the OCP for all classes except MAC, which has
  a routing of 1.0 for class 1 from the scheduler
  to URIRO.

4.5.4 Calculation of workload parameters

The workload parameters are those values which will vary from session to session, such as number of transfers to disc or tape, amount of OCP time used, number of jobs completing. The idea is not to calculate these parameters every time. During calibration they must obviously all be calculated. However, during validation only a subset of these workload parameters is recalculated. These are the values that typically define the size of the workload, such as concurrency.

The workload parameters and their methods of computation are as follows:

- OCP service rate ; referring to figure 4.4 it can be seen that from the OCP a job will branch one of 7 ways (i.e. to the UTERM, UVSI, UDISC, URIRO, UTAPE or UPAPER centres, or in the case of the MAC chain, it may depart from the network). It will then return to the OCP before taking one of these routings again. The total number of routings possible from the OCP is therefore:

    TOTAL = JOBS + VSIS + DISC TRANSFERS + RIROS +
            TAPE TRANSFERS + RECORDS SPOOLED

The time spent at the OCP each time round the loop can be calculated as follows:

    SERVICE TIME = OCP TIME / TOTAL

The OCP service rate is the inverse of this figure.

- branching probabilities from the OCP : the branching probabilities for each micro centre are calculated thus:

    UTERM = JOBS / TOTAL
    UVSI = VSIS / TOTAL

61

```
UDISC  = DISC TRANSFERS / TOTAL
URIRO  = RIROS / TOTAL
UTAPE  = TAPE TRANSFERS / TOTAL
UPAPER = RECORDS SPOOLED / TOTAL
```

The departure probability for the MAC chain is simply
RIROS / TOTAL. However, this does not need to be
explicitly inserted into the model as a parameter. SNAP
recognizes the fact that the branching probabilities do
not sum to one, and assumes the remaining probability
represents a departure from the network;

- branching probabilities from UDISC : from the UDISC
  centre jobs may branch to any of the disc drive centres.
  The probabilities for each branch are simply calculated
  as the number of transfers to the relevant disc drive
  divided by the sum of all these disc transfers;

- MAC arrival rate : this is determined by:

  ARRIVAL RATE = (RIROS * MPL) / ELAPSED TIME

- concurrencies : for the BATCH and SYSTEM workload chains
  the MPL is simply obtained from the concurrency report
  of the NRP (see also section 4.6.2).

4.6 MODEL CALIBRATION

4.6.1 The calibration process

Once the configuration and functional models have been designed
the calibration stage can begin. Calibration is the process
whereby it is determined whether or not the performance indices
predicted by the model fall within a certain acceptable error
margin when compared to the actual performance indices as
displayed by the real system. The calibration process is shown
in figure 4.8 below. One can approach the calibration phase in

```
┌──────────────┐      ┌──────────────┐      ┌──────────────┐
│   ACTUAL     │ ───> │   ACTUAL     │ ───> │   ACTUAL     │
│  WORKLOAD    │      │ CONFIGURATION│      │ PERFORMANCE  │
└──────────────┘      └──────────────┘      └──────────────┘
       │                                          ▲
       ▼                                          │
  MEASUREMENT                                 CALIBRATION
       │                                          │
       ▼                                          ▼
┌──────────────┐      ┌──────────────┐      ┌──────────────┐
│   WORKLOAD   │ ───> │CONFIGURATION │ ───> │  PREDICTED   │
│    MODEL     │      │    MODEL     │      │ PERFORMANCE  │
└──────────────┘      └──────────────┘      └──────────────┘
```
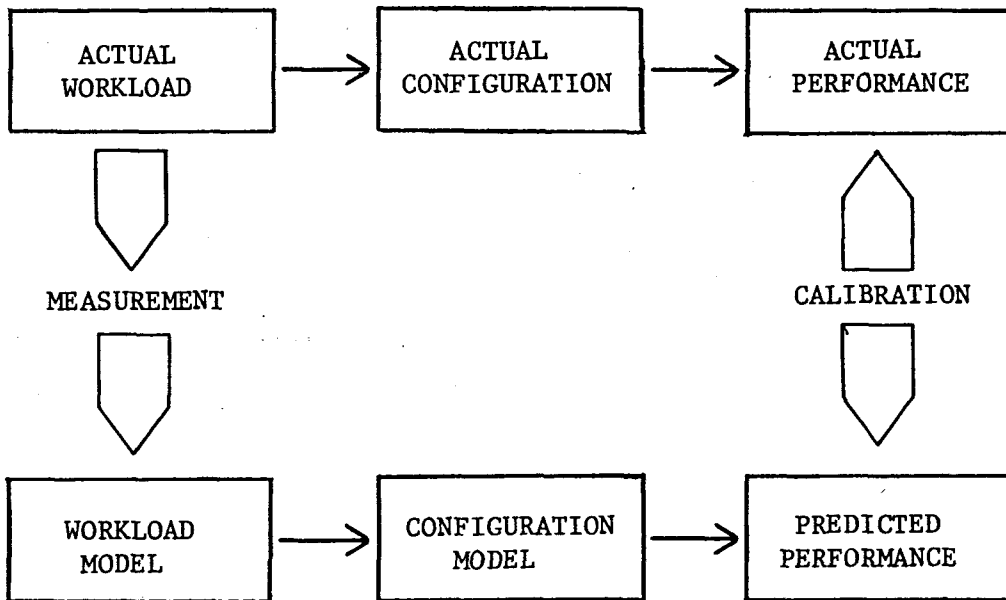
Figure 4.8   The calibration process

many ways. For this reason it is regarded as more of an art
than a science. Many iterations of the process are normally
required before the specified level of accuracy is achieved.
These iterations generally take the form of varying some
parameter that is not directly measureable until the indices
fall within the accepted error margin, or if one is fortunate
enough to have all the required parameters available, then the

63

model will need to be refined and/or the workload parameters will need to be calculated to a greater level of significance.

Failure of the model to achieve the required accuracy could be attributed to one of the following:

- model - incorrect design, incomplete design or an over simplified design;

- parameters - incorrect calculation of the workload parameters, level of significance of the workload parameters not great enough or over simplification or errors in the design of the functional model;

- solution method - incorrect or too simple method used to solve the network model. (Using SNAP, we could be fairly confident that this would not be the case!!).

The calibration process is generally performed using one set of data that is deemed to be representative of the workload. Having arrived at a model design and set of parameters that display the required accuracy, one then proceeds to the validation process, which tests the robustness of the model using performance data gathered from other monitored sessions.

4.6.2 GFSA model calibration

Calibration of the GFSA model was performed using data collected during the session of 4 June 1981 (session identity JUNO41). A series of sessions were inspected before deciding on JUNO41. This session was chosen because it lasted the full day and also it displayed no obviously abnormal behaviour. As we were only interested in the 'office hours' period, the data was run through the NRP with the time selection facility set to select only those jobs falling between 0858 - 1255 and 1400 - 1630. The NRP was run three times to select the three

64

workload types, BATCH, MAC and SYSTEM (see section 4.4 and
4.5.2.2). Appendix B shows the NRP listings for JUN041.

The workload parameters were calculated in the manner described
in section 4.5.2. The parameters for the three workload classes
are shown in tables 4.2 to 4.4. As mentioned earlier, the value
for the scheduler service rate is unknown and so this is the
parameter that we vary until the performance indices predicted
by the model agree with those calculated directly using the
NRP.

The model was first run with the scheduler service rate set to
zero. Unfortunately, SNAP can only run with an integral
concurrency value. The SNAP analysis therefore had to be run
four times varying the concurrency parameter values for the
BATCH and SYSTEM workloads as follows:

BATCH	SYSTEM
3	3
3	4
4	3
4	4

A four-way linear interpolation was then performed to obtain
the performance indices in terms of 3.35 BATCH jobs and 3.79
SYSTEM jobs.

It is true that the SYSTEM workload consists of jobs that run
for the whole session and so by including these jobs we are
including work that was performed outside the 0800 - 1630
timeslot. However, it must be remembered that we are dealing
with averages and branching probabilities. So long as the same
type of work is done, eg. the spooler jobs access the discs in
the same proportion for their print files throughout the
session, then the figures will not be too incorrect. This is
acceptable because, firstly, the spooler jobs and master

65

OCP SERVICE RATE   =   91.8628 per second

MAC ARRIVAL RATE   =   0.3857 per second

BRANCHING PROBABILITIES

FROM	TO	VISITS	PROBABILITY
	UTERM	162	0.0004
	UVSI	139177	0.3376
OCP	UDISC	255858	0.6207
	URIRO	8622	0.0209
	UTAPE	8391	0.0204
		412210	1.0000
	SYS002	26519	0.1038
	SYS003	8089	0.0317
	SYS001	166508	0.6509
	GFSA13	13148	0.0515
	ED16	0	0.0
UDISC	ED17	63	0.0002
	GFSA01	5520	0.0216
	GFSA02	19532	0.0764
	GFSA03	11520	0.0451
	GFSA12	4407	0.0172
	SYS004	416	0.0016
		255722	1.0000

Table 4.2   Workload parameters for the MACRESPOND workload

```
        OCP SERVICE RATE     = 83.4481 per second
        AVERAGE CONCURRENCY  =  3.35   jobs



               BRANCHING PROBABILITIES


   FROM            TO          VISITS     PROBABILITY


              ┌── UTERM          62         0.0002
              ├── UVSI         38330        0.1349
   OCP ───────┼── UDISC       210285        0.7402
              ├── URIRO         132         0.0004
              └── UTAPE        35319        0.1243

                              ──────       ──────
                              284128        1.0000



              ┌── SYS002       13080        0.0622
              ├── SYS003       22823        0.1085
              ├── SYS001       97084        0.4618
              ├── GFSA13        5283        0.0251
              ├── ED16            0         0.0
   UDISC ─────┼── ED17            0         0.0
              ├── GFSA01        5805        0.0276
              ├── GFSA02       22009        0.1047
              ├── GFSA03        1335        0.0063
              ├── GFSA12       41201        0.1959
              └── SYS004        1665        0.0079

                              ──────       ──────
                              210285        1.0000
```

Table 4.3 Workload parameters for the BATCH workload

```
           OCP SERVICE RATE    = 543.7725 per second
        AVERAGE CONCURRENCY =    3.79   jobs



                 BRANCHING PROBABILITIES


    FROM            TO        VISITS      PROBABILITY


                 ┌── UTERM         5        0.00001
                 ├── UVSI      59527        0.0694
    OCP ─────────┼── UDISC     86082        0.1004
                 ├── UTAPE       306        0.0004
                 └── UPAPER   711578        0.82979

                             ─────────   ──────────
                             857498        1.00000



                 ┌── SYS002     4340        0.0504
                 ├── SYS003    14925        0.1734
                 ├── SYS001    39063        0.4537
                 ├── GFSA13       30        0.0004
                 ├── ED16       6173        0.0717
    UDISC ───────┼── ED17       4082        0.0474
                 ├── GFSA01      151        0.0017
                 ├── GFSA02     8089        0.0940
                 ├── GFSA03     3786        0.0440
                 ├── GFSA12     5241        0.0609
                 └── SYS004      202        0.0024

                             ─────────   ──────────
                             86082         1.0000
```

Table 4.4   Workload parameters for the SYSTEM workload

operator job do perform work of the same nature throughout the day, and secondly, their resource demands are so small when compared with the BATCH and MAC workloads that an inaccuracy in the SYSTEM workload will not affect the model as a whole very significantly.

The results for this first run are shown in tables 4.5 to 4.7, where they are compared with the measured results. These 'measured results' are obtained either directly from the NRP (eg. OCP utilization) or by performing some calculation on the figures obtained from the NRP (eg. each disc utilization is obtained by the following calculation :

number of accesses X access time (= 0.03833 seconds)

elapsed time for this class

The percentage errors presented in table 4.5 show that the MAC workload class is excellently represented by the model. This is due to the fact that the MAC workload is modelled as an open chain. The amount of work that can be done is restricted by the arrival rate. The scheduler service rate does not affect the performance indices. This is because there is one visit to the scheduler for each interaction. Changing the time spent at the scheduler only serves to change the response time for this open chain. The abnormal error for disc drive ED17 is simply because the branching probability was not calculated to a sufficient degree of accuracy. Very few transfers take place to ED17 and so the associated routing probability is very small. The significance was increased by one more decimal place. The result is apparant in table 4.8.

The BATCH and SYSTEM classes are modelled as closed chains and thus behave differently. Tables 4.6 and 4.7 are encouraging in that the ₋percentage errors are of the same magnitude for each centre. It therefore seems likely that adding in the scheduler delay will cause these percentage errors to decrease at the

69

PERFORMANCE INDICE	RESULTS		% ERROR
	MEASURED	MODEL	
UTILIZATIONS (%):			
OCP	20.07	20.0889	+0.094
SYS002	16.4813	16.5541	+0.44
SYS003	13.3208	13.3777	+0.43
SYS001	28.5541	28.6761	+0.43
GFSA13	2.2547	2.2689	+0.63
ED16	0.0	0.0	-
ED17	0.0108	0.0088	-22.73
GFSA01	0.9466	0.9516	+0.53
GFSA02	3.3495	3.3659	+0.49
GFSA03	1.9756	1.9869	+0.57
GFSA12	0.7557	0.7578	+0.28
SYS004	1.5499	1.5540	+0.26
DC10	2.8575	2.8601	+0.091
DC40	22.2018	22.2014	-0.002
MC00	0.2802	0.2809	+0.25
MC10	0.2802	0.2809	+0.25
THROUGHPUT RATES: (per second)			
UTERM	0.00725	0.0074	+2.07
UVSI	6.2268	6.2302	+0.055
UDISC	11.4471	11.4546	+0.066
URIRO	0.3857	0.3857	0.0
UTAPE	0.3754	0.3765	+0.29

Table 4.5   Results of the first calibration run   (scheduler service rates set to infinity) for the MACRESPOND workload class

PERFORMANCE INDICE	RESULTS		% ERROR
	MEASURED	MODEL	
UTILIZATIONS (%):			
OCP	16.51	41.1520	+149.0
SYS002	5.9820	14.9896	+150.6
SYS003	7.7895	19.5161	+150.5
SYS001	18.0107	45.1478	+150.6
GFSA13	0.9801	2.4539	+150.4
ED16	0.0	0.0	–
ED17	0.0	0.0	–
GFSA01	1.0769	2.6983	+150.6
GFSA02	4.0830	10.2359	+150.7
GFSA03	0.2477	0.6159	+148.6
GFSA12	7.6435	19.1521	+150.6
SYS004	0.3334	0.8252	+147.5
DC10	1.7337	4.3302	+149.8
DC40	1.7087	4.1084	+140.4
MC00	1.2757	3.1855	+149.7
MC10	1.2757	3.1855	+149.7
THROUGHPUT RATES: (per second)			
UTERM	0.0030	0.0069	+130.0
UVSI	1.8552	4.6325	+149.7
UDISC	10.1777	25.4189	+149.8
URIRO	0.0064	0.0137	+114.1
UTAPE	1.7094	4.2685	+149.7

Table 4.6   Results of the first calibration run (scheduler service
rates set to infinity) for the BATCH workload class

| PERFORMANCE INDICE | RESULTS | | % ERROR |
	MEASURED	MODEL	
UTILIZATIONS (%)			
OCP	2.72	7.7397	+184.5
SYS002	2.0998	6.4359	+206.5
SYS003	2.7515	8.4349	+206.6
SYS001	2.4051	7.3734	+206.6
GFSA13	0.00185	0.0065	+251.4
ED16	0.3801	1.1653	+206.6
ED17	0.2513	0.7703	+206.5
GFSA01	0.0093	0.0276	+196.8
GFSA02	0.4980	1.5277	+206.8
GFSA03	0.2331	0.7151	+206.8
GFSA12	0.3227	0.9897	+206.7
SYS004	0.0124	0.0390	+214.5
DC10	0.3748	1.1452	+205.5
DC40	0.2056	0.6281	+205.5
MC00	0.0037	0.0126	+240.5
MC10	0.0037	0.0126	+240.5
THROUGHPUT RATES: (per second)			
UTERM	0.00008	0.0004	+400.0
UVSI	0.9562	2.9208	+205.5
UDISC	1.3827	4.2255	+205.6
UTAPE	0.0049	0.0168	+242.9
UPAPER	11.4302	34.9227	+205.5

Table 4.7   Results of the first calibration run (scheduler service
rates set to infinity) for the SYSTEM workload class

same rate.

Calibration now began in earnest. The accuracy of the MAC chain was considered as being more than adequate. Only one change was made to it, namely the scheduler service rate was set to 2 per second. This was designed to add a scheduling delay to the MAC chain so as to produce a more realistic response time figure. This value is more intuitively than scientifically based as there is no monitored figure available.

Two changes were made to the BATCH chain:

- the branching probabilities from the OCP to the UTERM and URIRO centres were given a greater degree of significance;

- the scheduler service rate was varied. It was finally set at 0.0039 per second.

Turning to the SYSTEM chain, the UTERM and UTAPE branching probabilities were given added significance. The scheduler service rate was varied but without much success. This is because the scheduler is only visited at the end of a job. There were only 5 SYSTEM jobs, so this results in a very small routing probability to the scheduler. It was decided to calibrate the SYSTEM workload on the printer service rate. This produced an immediate improvement. The service rate was finally set at 6.4 per second. (The SYSTEM scheduler service rate was set at 0.1 per second).

The final calibration results are presented in tables 4.8 to 4.10. The terminal response time is 4.98 seconds.

An example of a SNAP output listing is given in appendix E.

PERFORMANCE INDICE	RESULTS		% ERROR
	MEASURED	MODEL	
UTILIZATIONS (%):			
OCP	20.07	20.0889	+0.094
SYSOO2	16.4813	16.5541	+0.44
SYSOO3	13.3208	13.3777	+0.43
SYSOO1	28.5541	28.6761	+0.43
GFSA13	2.2547	2.2667	+0.53
ED16	0.0	0.0	-
ED17	0.0108	0.0110	+1.85
GFSAO1	0.9466	0.9516	+0.53
GFSAO2	3.3495	3.3659	+0.49
GFSAO3	1.9756	1.9869	+0.57
GFSA12	0.7557	0.7578	+0.28
SYSOO4	1.5499	1.5540	+0.26
DC10	2.8575	2.8601	+0.091
DC40	22.2018	22.2014	-0.002
MCOO	0.2802	0.2809	+0.25
MC10	0.2802	0.2809	+0.25
THROUGHPUT RATES: **(per second)**			
UTERM	0.00725	0.0074	+2.07
UVSI	6.2268	6.2302	+0.055
UDISC	11.4471	11.4546	+0.066
URIRO	0.3857	0.3857	0.0
UTAPE	0.3754	0.3765	+0.29

Table 4.8 Results of the final calibration run (scheduler service
rates set to infinity) for the MACRESPOND workload class

| PERFORMANCE INDICE | RESULTS | | % ERROR |
	MEASURED	MODEL	
UTILIZATIONS (%)			
OCP	16.51	16.3928	-0.71
SYS002	5.9820	5.9708	-0.19
SYS003	7.7895	7.7738	-0.20
SYS001	18.0107	17.9825	-0.16
GFSA13	0.9801	0.9774	-0.28
ED16	0.0	0.0	-
ED17	0.0	0.0	-
GFSA01	1.0769	1.0748	-0.20
GFSA02	4.0830	4.0771	-0.14
GFSA03	0.2477	0.2453	-0.97
GFSA12	7.6453	7.6284	-0.20
SYS004	0.3334	0.3318	-0.48
DC10	1.7337	1.7248	-0.51
DC40	1.7087	1.6951	-0.80
MC00	1.2757	1.2689	-0.53
MC10	1.2757	1.2689	-0.53
THROUGHPUT RATES:			
(per second)			
UTERM	0.0030	0.0030	-0.0
UVSI	1.8552	1.8454	-0.53
UDISC	10.1777	10.1245	-0.52
URIRO	0.0064	0.0063	-1.56
UTAPE	1.7094	1.7004	-0.53

Table 4.9   Results of the final calibration run (scheduler service
rates set to infinity) for the BATCH workload class

PERFORMANCE INDICE	RESULTS		% ERROR
	MEASURED	MODEL	
UTILIZATIONS (%)			
OCP	2.72	2.77	+1.84
SYS002	2.0998	2.1725	+3.46
SYS003	2.7515	2.8011	+1.80
SYS001	2.4051	2.4347	+1.23
GFSA13	0.00185	0.00189	+2.16
ED16	0.3801	0.3872	+1.87
ED17	0.2513	0.2580	+2.67
GFSA01	0.0093	0.0095	+2.15
GFSA02	0.4980	0.5011	+0.62
GFSA03	0.2331	0.2367	+1.54
GFSA12	0.3227	0.3285	+1.80
SYS004	0.0124	0.0127	+2.42
DC10	0.3748	0.3810	+1.65
DC40	0.2056	0.2092	+1.75
MC00	0.0037	0.0038	+2.70
MC10	0.0037	0.0038	+2.70
THROUGHPUT RATES:			
(per second)			
UTERM	0.00008	0.0001	-20.00
UVSI	0.9562	0.9725	+1.70
UDISC	1.3827	1.4062	+1.70
UTAPE	0.0049	0.0050	+2.04
UPAPER	11.4302	11.6245	+1.70

Table 4.10 Results of the final calibration run (scheduler service rates set to infinity) for the SYSTEM workload class

4.7 MODEL VALIDATION

4.7.1 The validation process

Once the model has been calibrated to a specified degree of accuracy, it needs to be validated. We need to be confident that when we come to use the model for prediction purposes, it will produce results that lie within a known error margin. Validation implies taking the workload parameters calculated from another session, changing only those parameters that have a direct bearing on the size of the workload, and seeing whether the performance indices now given by the model reflect the same or nearly the same degree of accuracy as that given during the calibration stage. If they do not, it does not necessarily signify failure. One needs to determine the reasons behind any major deviation from the measured values. Such inaccuracies may only be present in certain areas of the model. Valid predictions could possibly still be made for the other components. It is shown below, for example, that the disc traffic is very erratic on the GFSA system. However, the validation results obtained for the OCP utilizations are very good. One can therefore perform OCP utilization predictions with a certain amount of confidence.

4.7.2 GFSA model validation

Validation of the GFSA model was performed using two sessions, namely JUN081 and JUN182. These sessions were chosen firstly because they appeared to display no abnormalities, and secondly because their workloads differed somewhat from JUN041.

JUN081 was a morning session lasting for just under five hours. The concurrencies for the BATCH and SYSTEM workloads were 2.52 and 2.84 jobs respectively. The arrival rate for the MACRESPOND workload was 0.2249 per second. This session therefore had a

considerably lighter workload than JUN041, indicating that it would be a good test of the model validity.

JUN182 lasted for a full day. The BATCH concurrency was 4.23 jobs (higher than JUN041), whilst the SYSTEM concurrency was 2.89 jobs and the MACRESPOND arrival rate was 0.2989 (both lower than the values for JUN041).

The validation process requires that only those parameters that have a bearing on the size of the workload be changed. If the model were perfectly behaved, this would imply changing only the concurrency values for the closed chains and the arrival rate for the open chain. It was soon realized that the GFSA workload was not very well behaved, and that further changes had to be made. In the end, three runs were performed on each validation session. These are termed Model 1, 2 and 3 and contain the following variations:

- Model 1 : BATCH and SYSTEM concurrencies and MACRESPOND arrival rate changed to reflect the measured values;

- Model 2 : OCP service rates altered for each class;

- Model 3 : branching probabilities from the OCP to each micro server changed for each class.

The two validation sessions and their associated validation results are discussed below.

4.7.2.1 Validation session JUN182

The workload parameters for the three workload groups are shown in tables 4.11 to 4.13. These values can be compared with those given in tables 4.2 to 4.4 for session JUN041.

```
OCP SERVICE RATE   =   94.6225 per second
MAC ARRIVAL RATE   =    0.2989 per second
```

BRANCHING PROBABILITIES

FROM	TO	VISITS	PROBABILITY
	UTERM	181	0.00034
	UVSI	228389	0.43218
OCP	UDISC	273383	0.51734
	URIRO	7718	0.01460
	UTAPE	18783	0.03554
		528454	1.00000
	SYS002	7574	0.0277
	SYS003	4352	0.0159
	SYS001	189360	0.6923
	GFSA13	15349	0.0561
	ED16	0	0.0
UDISC	ED17	0	0.0
	GFSA01	10695	0.0391
	GFSA02	19805	0.0724
	GFSA03	6310	0.0231
	GFSA12	18833	0.0689
	SYS004	945	0.0035
		273223	1.0000

Table 4.11   MACRESPOND workload parameters for session JUN182

79

```
              OCP SERVICE RATE      =   87.6727 per second
         AVERAGE CONCURRENCY =     4.23   jobs



                  BRANCHING PROBABILITIES


    FROM              TO           VISITS      PROBABILITY


                  ┌── UTERM          70        0.00009
                  ├── UVSI        68063        0.08821
    OCP ──────────┼── UDISC      650279        0.84275
                  ├── URIRO         44        0.00006
                  └── UTAPE       53154        0.06889

                                 ─────────      ─────────

                                 771610        1.00000




                  ┌── SYS002       25806        0.0397
                  ├── SYS003        3946        0.0061
                  ├── SYS001      523585        0.8051
                  ├── GFSA13        5977        0.0092
                  ├── ED16             0        0.0
    UDISC ────────┼── ED17         27317        0.0420
                  ├── GFSA01       18125        0.0279
                  ├── GFSA02       31079        0.0478
                  ├── GFSA03       10475        0.0161
                  ├── GFSA12        3771        0.0058
                  └── SYS004         198        0.0003

                                 ─────────      ─────────

                                 650279        1.0000
```

Table 4.12   BATCH workload parameters for session JUN182

80

OCP SERVICE RATE   = 535.8413 per second

AVERAGE CONCURRENCY =   2.89   jobs

BRANCHING PROBABILITIES

FROM	TO	VISITS	PROBABILITY
	UTERM	5	0.00001
	UVSI	56839	0.10313
OCP	UDISC	66839	0.12127
	UTAPE	268	0.00049
	UPAPER	427207	0.77511
		551158	1.00000
	SYS002	4249	0.0636
	SYS003	9331	0.1396
	SYS001	32182	0.4815
	GFSA13	18	0.0003
	ED16	5724	0.0856
UDISC	ED17	5618	0.0840
	GFSA01	2656	0.0397
	GFSA02	2772	0.0415
	GFSA03	113	0.0017
	GFSA12	3784	0.0566
	SYS004	392	0.0059
		66839	1.0000

Table 4.13   SYSTEM workload parameters for session JUN182

The first is that the OCP service rates compare exceptionally well. It can therefore be expected that only a slight difference will be exhibited between model 1 and model 2. This is indeed the case. (The results for each class are given in tables 4.14 to 4.16).

The results for the first two models are, however, rather erratic. If the percentage errors were all of the same magnitude, the deviation from the measured value could then be attributed to the scheduler service rate setting. The fact that they are not of the same order is due to the varying behaviour of the GFSA workload. Looking at the branching probabilities from the OCP to the micro servers, it is apparent that significant differences exist between the corresponding figures for JUN041 and JUN182. A number of sessions should be analyzed in an attempt to determine mean values for these routing probabilities. One needs to be aware of the deviation that is possible, so that predictions can be performed using a three-way analysis strategy, namely a worst, average and optimum case analysis.

These branching probabilities are changed in model 3 to the correct values for JUN182. This may at first appear to be an illegal move for the validation process. However, the aim is to verify that the scheduler service rate values, reached during the calibration phase, are valid. If satisfactory results are achieved in this way, one can then be confident that the model structure is an accurate representation of the system. Major variations in the workload from day to day can not be resolved within the model itself. This does not mean that the model can not be shown to be valid. GFSA is unfortunate in having such a varying workload, but meaningful evaluations can still be performed. One just needs to be aware of the constraints of the system.

PERFORMANCE INDICE	RESULTS						
	MEASURED	MODEL 1	% ERROR	MODEL 2	% ERROR	MODEL 3	% ERROR
UTILIZATIONS (%) :							
OCP	21.43	15.57	−27.34	15.11	−29.49	21.64	0.98
DC10	3.1060	2.2164	−28.64	2.2164	−28.64	3.0275	−2.53
DC40	17.7333	17.2051	−2.98	17.2051	−2.98	17.8168	0.47
MC	0.5428	0.2177	−59.89	0.2177	−59.89	0.5430	0.04
THROUGHPUT RATES :							
(per second)							
UTERM	0.0070	0.0057	−18.57	0.0057	−18.57	0.0070	0.0
UVSI	8.8444	4.8281	−45.41	4.8281	−45.41	8.8476	0.04
UDISC	10.5868	8.8768	−16.15	8.8768	−16.15	10.5910	0.04
URIRO	0.2989	0.2989	0.0	0.2989	0.0	0.2989	0.0
UTAPE	0.7274	0.2917	−59.90	0.2917	−59.90	0.7276	0.03

Table 4.14   MACRESPOND validation results for session JÚN182

PERFORMANCE INDICE	RESULTS						
	MEASURED	MODEL 1	% ERROR	MODEL 2	% ERROR	MODEL 3	% ERROR
UTILIZATIONS (%) :							
OCP	34.27	20.82	-39.25	19.91	-41.90	36.27	5.84
DC10	5.6490	2.1905	-61.22	2.2005	-61.05	4.3070	-23.75
DC40	1.2291	1.5817	28.69	1.5889	29.27	3.0383	147.20
MC	1.5911	1.6116	1.29	1.6189	1.75	1.6347	2.74
THROUGHPUT RATES : (per second)							
UTERM	0.00281	0.00381	35.59	0.00381	35.59	0.00289	2.85
UVSI	2.7302	2.3437	-14.16	2.3544	-13.76	2.8048	2.73
UDISC	26.0842	12.8583	-50.70	12.9171	-50.48	26.7969	2.73
URIRO	0.00176	0.00799	353.98	0.00801	355.11	0.00188	6.82
UTAPE	2.1321	2.1595	1.29	2.1694	1.75	2.1905	2.74

Table 4.15   BATCH validation results for session JUN182

PERFORMANCE INDICE	RESULTS						
	MEASURED	MODEL 1	% ERROR	MODEL 2	% ERROR	MODEL 3	% ERROR
UTILIZATIONS (%) :							
OCP	2.72	2.43	7.05	2.46	8.37	2.56	12.78
DC10	0.4436	0.3591	−19.05	0.3591	−19.05	0.4846	9.24
DC40	0.2313	0.1969	−14.87	0.1969	−14.87	0.2807	21.36
MC	0.0044	0.0035	−20.45	0.0035	−20.45	0.0050	13.64
THROUGHPUT RATES : (per second)							
UTERM	0.0001	0.0001	0.0	0.0001	0.0	0.0001	0.0
UVSI	1.2500	0.9157	−26.74	0.9157	−26.74	1.4174	13.39
UDISC	1.4700	1.3248	−9.88	1.3248	−9.88	1.6667	13.40
UTAPE	0.0059	0.0047	−20.34	0.0047	−20.34	0.0067	13.56
UPAPER	9.3952	10.9501	16.55	10.9503	16.55	10.6529	13.39

Table 4.16   SYSTEM validation results for session JUN 82

The validation results for model 3 are exceptional. ITR's past experience has shown that a percentage error of up to 30% is still acceptable. Tables 4.14 to 4.16 show errors of 0%, 2.7% and 13.5% for the three groups. Two points should be mentioned. The discrepancy in the BATCH URIRO rate is most likely due to the level of significance assigned to the branching probabilities. They were all increased to five decimal places for the validation run, but this is still not really enough for the centres with very low utilizations.

The second point concerns the accuracy of the DFC utilizations. These values are obviously dependent on the amount of traffic going to each disc drive. The behaviour of the branching probabilities from UDISC to the disc drives is even more erratic than those from the OCP. For example, ED17 caters for removable disc packs. On 4 June no BATCH job accessed this drive, while on 18 June over 4% of the total BATCH disc traffic was routed there. These probabilities can vary an enormous amount, and hence the disc utilizations are not reported in the validation results. The DFC utilizations will vary accordingly (note that DC40 also contains the RIRO traffic, hence the large value for the MACRESPOND workload group).

The validation values given for JUN182 show the model to be well designed and calibrated. Obviously one can not rely on the figures obtained from one validation session. A new session was therefore chosen, namely that of JUN081.

4.7.2.2 Validation session JUN081

The workload parameters for this session are shown in tables 4.17 to 4.19, while the validation results are

given in tables 4.20 to 4.22.

The same considerations apply for this session. The OCP
service rate for the MACRESPOND group is very close to
the JUNO41 value. However, the figure for the BATCH and
SYSTEM classes are very different to the corresponding
JUNO41 values. This is reflected in the validation
results for models 1 and 2.

The correction of the OCP branching probabilities once
again provides model 3 with excellent results.

The model has thus been shown to be remarkably accurate. One
final figure of interest is the MACRESPOND response time. The
value for JUN182 is 8.42 seconds and that for JUNO81 is 6.60
seconds.

```
        OCP SERVICE RATE   =   93.3661 per second
        MAC ARRIVAL RATE   =    0.2249 per second
```

BRANCHING PROBABILITIES

FROM	TO	VISITS	PROBABILITY
	UTERM	95	0.00033
	UVSI	106962	0.37536
OCP	UDISC	165638	0.58129
	URIRO	3899	0.01368
	UTAPE	8362	0.02934
		284956	1.00000
	SYS002	3750	0.0226
	SYS003	9605	0.0580
	SYS001	120669	0.7286
	GFSA13	7563	0.0456
	ED16	0	0.0
UDISC	ED17	22	0.0001
	GFSA01	1529	0.0092
	GFSA02	12509	0.0756
	GFSA03	8333	0.0503
	GFSA12	68	0.0004
	SYS004	1590	0.0096
		165638	1.0000

Table 4.17   MACRESPOND workload parameters for session JUN081

```
         OCP SERVICE RATE    = 102.167 per second
         AVERAGE CONCURRENCY =   2.52  jobs



                    BRANCHING PROBABILITIES


     FROM            TO         VISITS     PROBABILITY


                  ┌── UTERM        53        0.00020
                  ├── UVSI      39077        0.41394
     OCP ─────────┼── UDISC    215446        0.79361
                  ├── URIRO       46         0.00017
                  └── UTAPE     16954        0.06208

                              ─────────    ──────────
                               271476        1.00000


                  ┌── SYS002    17638        0.0820
                  ├── SYS003    22153        0.1028
                  ├── SYS001   113308        0.5259
                  ├── GFSA13     4598        0.0213
                  ├── ED16         84        0.0004
     UDISC ───────┼── ED17          0        0.0
                  ├── GFSA01     5421        0.0252
                  ├── GFSA02    11562        0.0536
                  ├── GFSA03    37867        0.1757
                  ├── GFSA12     2815        0.0131
                  └── SYS004        0        0.0

                              ─────────    ──────────
                               215446        1.0000
```

Table 4.18   BATCH workload parameters for session JUN081

```
OCP SERVICE RATE     = 312.3149 persecond
AVERAGE CONCURRENCY  =   2.84   jobs
```

BRANCHING PROBABILITIES

FROM	TO	VISITS	PROBABILITY
	UTERM	5	0.00002
	UVSI	27629	0.08540
OCP	UDISC	40585	0.12545
	UTAPE	0	0.0
	UPAPER	255291	0.78913
		323510	1.00000
	SYS002	1745	0.0430
	SYS003	15719	0.3873
	SYS001	20803	0.5125
	GFSA13	77	0.0019
	ED16	0	0.0
UDISC	ED17	0	0.0
	GFSA01	321	0.0079
	GFSA02	949	0.0234
	GFSA03	640	0.0158
	GFSA12	0	0.0
	SYS004	331	0.0082
		40585	1.0000

Table 4.19   SYSTEM workload parameters for session JUN081

PERFORMANCE INDICE	RESULTS						
	MEASURED	MODEL 1	% ERROR	MODEL 2	% ERROR	MODEL 3	% ERROR
UTILIZATIONS (%) :							
OCP	17.61	11.96	-32.08	11.77	-33.16	17.85	1.36
DC10	2.7384	1.6677	-39.10	1.6677	-39.10	2.5068	-8.46
DC40	16.0038	12.9455	-19.11	12.9455	-19.11	13.4501	-15.96
MC	0.3599	0.1638	-54.49	0.1638	-54.49	0.3600	0.03
THROUGHPUT RATES : (per second)							
UTERM	0.0054	0.0043	20.37	0.0043	20.37	0.0054	0.0
UVSI	6.1692	3.6328	-41.11	3.6328	-41.11	6.1708	0.03
UDISC	9.5535	6.6791	-30.09	6.6791	-30.09	9.5562	0.03
URIRO	0.2249	0.2249	0.0	0.2249	0.0	0.2249	0.0
UTAPE	0.4823	0.2195	-54.49	0.2195	-54.49	0.4823	0.0

Table 4.20   MACRESPOND validation results for session JUNO81

PERFORMANCE INDICE	RESULTS						
	MEASURED	MODEL 1	% ERROR	MODEL 2	% ERROR	MODEL 3	% ERROR
UTILIZATIONS (%) :							
OCP	17.11	12.83	-25.01	10.61	-38.00	15.21	-11.10
DC10	2.5453	1.3504	-46.95	1.3664	-46.52	2.1001	-17.50
DC40	1.4254	0.9751	-31.60	0.9867	-30.78	1.5161	6.36
MC	0.8080	0.9935	22.96	1.0053	24.42	0.7201	-10.88
THROUGHPUT RATES : (per second)							
UTERM	0.0034	0.0024	29.41	0.0024	29.41	0.0031	-8.82
UVSI	2.5102	1.4448	-42.44	1.4620	-41.76	2.2374	-10.87
UDISC	13.8400	7.9269	-42.72	8.0209	-42.05	12.3359	-10.87
URIRO	0.0030	0.0049	63.33	0.0050	66.67	0.0027	-10.00
UTAPE	1.0827	1.3313	22.96	1.3471	24.42	0.9650	-10.87

Table 4.21   BATCH validation results for session JUN081

PERFORMANCE INDICE	RESULTS						
	MEASURED	MODEL 1	% ERROR	MODEL 2	% ERROR	MODEL 3	% ERROR
UTILIZATIONS (%) :							
OCP	4.47	2.41	−46.09	4.18	−6.49	4.33	−3.12
DC10	0.5162	0.3571	−30.82	0.3554	−31.15	0.4581	−11.26
DC40	0.1823	0.1959	7.46	0.1949	6.91	0.2502	37.25
MC	0.0	0.0035	−	0.0035	−	0.0	0.0
THROUGHPUT RATES : (per second)							
UTERM	0.0002	0.0001	−50.0	0.0001	−50.0	0.0002	0.0
UVSI	1.1401	0.9108	−20.11	0.9063	−20.51	1.1562	1.4122
UDISC	1.6747	1.3176	−21.32	1.3112	−21.71	1.6984	1.4152
UTAPE	0.0	0.0047	−	0.0047	−	0.0	0.0
UPAPER	10.5340	10.8905	3.3843	10.8375	2.881	10.6836	1.4202

Table 4.22   SYSTEM validation results for session JUN081

5. PERFORMANCE PREDICTIONS

5.1 INTRODUCTION

The validated model can now be used to answer what are commonly termed "what-if" questions, i.e. What will happen to the machine performance If a certain change is made. The envisaged changes may apply to the workload or the system configuration, and may cause the performance to increase or decrease.

Predictions can be divided into three categories. These are represented by the diagrams in figures 5.1, 5.2 and 5.3. In each diagram the user population U supplies a workload W_1 to a system configuration C_1 giving a measured performance P_{11}. The first
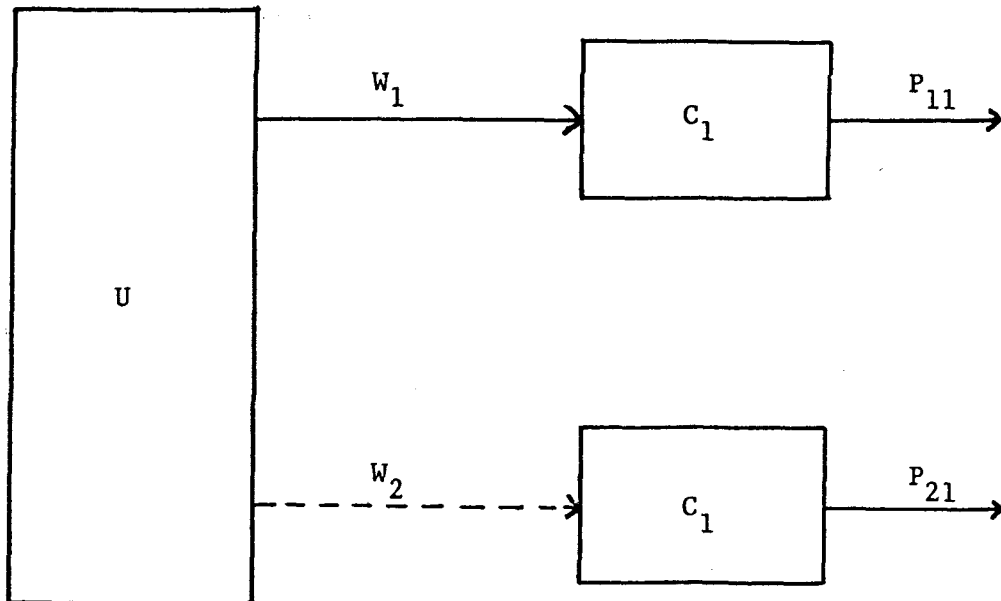
Figure 5.1   Performance prediction methodology 1

prediction methodology, shown in figure 5.1, is to increase the workload to a value W_2, apply this to the present configuration C_1 and observe the predicted performance P_{21} given by the model.

In figure 5.2 the workload is left unchanged. However, the configuration is altered in some way to produce a configuration C_2. When the model is run with the new parameter values it produces a predicted performance P_{12}.
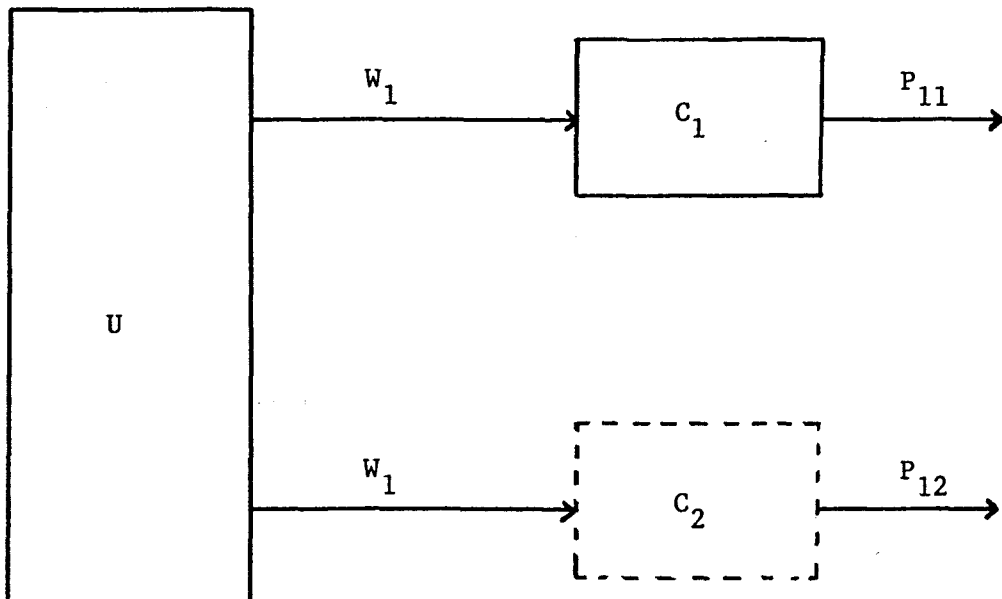
Figure 5.2   Performance prediction methodology 2

The final prediction methodology involves changing both the workload and the system configuration. Figure 5.3 shows how workload W_2, when applied to configuration C_2, produces a predicted performance P_{22}.
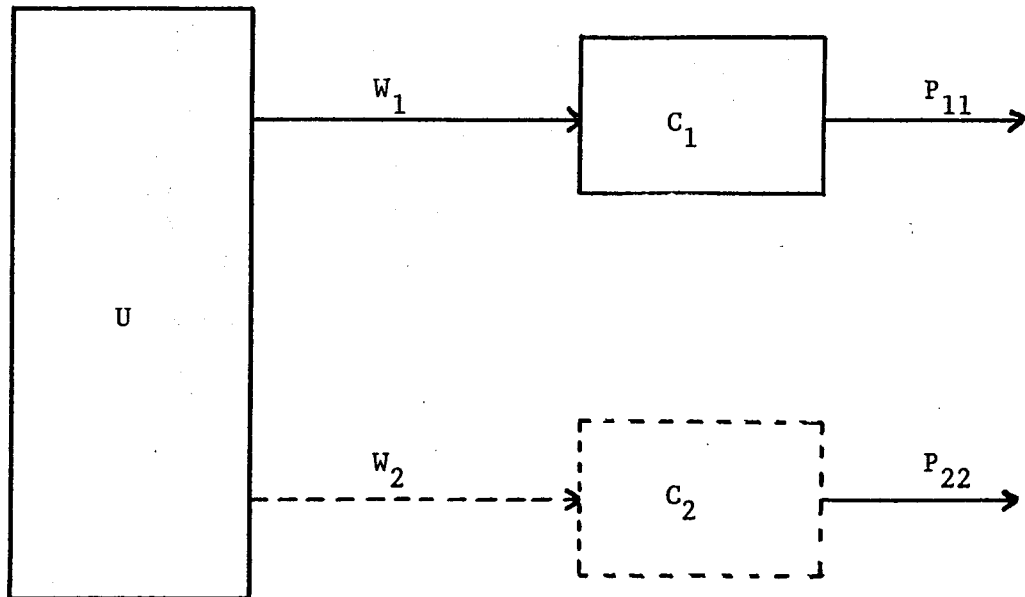
Figure 5.3   Performance prediction methodology 3

The recognition of these three methodologies provides an insight into the manner in which predictions should be carried out. There are obviously an enormous number of possible changes one can make to the workload and to the configuration. It is essential that meaningful predictions be made. Otherwise, unnecessary processing power will be consumed, as well as a lot of time being spent in analyzing the results.

The following section provides a number of examples which were performed for GFSA management. Idealy, the "what-if" questions should be posed by management. However, with a non-technical management structure, it falls to the performance analyst to provide not only the answers, but also the questions. Only a very broad directive can be expected from management.

5.2 PERFORMANCE PREDICTION EXAMPLES

The GFSA user community has only one performance concern, and that is the terminal response time. Terminal usage is fairly undisciplined. Users are encouraged to submit long program runs and compilations via a background profile, but this is not compulsory.

This strategy does not ease the performance analyst's task. There is, however, one advantage. Batch turnaround time is of no concern to any of the users. This means that performance tuning does not have to be a trade-off between batch throughput and terminal response time. In the examples that follow, only the response time and various centre utilizations will be used as performance indicators.

Looking at the final calibration results in tables 4.8 to 4.10, it can be seen that disc SYS001 has a much higher utilization than any other disc. SYS002 and SYS003 also have relatively high usage. The disc controller DC40 is heavily utilized by the MACRESPOND workload group. This is due to the RIRO traffic. It therefore seems likely that the disc subsystem could be causing a bottleneck situation. It is less likely to be caused by a lack of OCP processing power.

The examples shown below are divided into two sections. The first analyses the system as per the calibrated model of JUN041. The second performs the same set of analyses, but with the disc traffic equalized to all the disc drives (i.e. the branching probabilities from UDISC are equalized). In reality one would never be able to equalize all the traffic. However, the analysis is important in showing just how bad the bottleneck situation is.

Three prediction runs are performed for each configuration:

- the number of terminals is increased;

- the number of batch jobs is increased;
- the OCP speed (rated in MIPS - millions of instructions per second) is increased.

Three performance indicators are used. These are:

- the OCP utilization for each workload group, as well as the total OCP utilization;
- the utilization of the service centres representing SYS001, DC10 and DC40;
- the terminal (MACRESPOND) response time.

A graph is plotted for each set of results. There are thus nine plots for each section. All the plots have been drawn using a BATCH concurrency of 3 and a SYSTEM concurrency of 4, in order to avoid the interpolation overhead.

5.2.1 JUN041 configuration

Figures 5.1 to 5.3 show the results obtained when increasing the number of terminals (this is done by increasing the arrival rate proportionately. GFSA currently has 16 terminals. The assumption is made that one more terminal will lead to 1/16 more work being performed, two more terminals will give 2/16 more work, and so on).

The MACRESPOND OCP utilization increases steadily, whilst the BATCH OCP utilization decreases slightly. For the SYSTEM class it remains virtuall unchanged.

Steady increases are shown in the utilization of SYS001, DC10 and DC40. Note that the utilization of DC40 increases faster than that of DC10, due to the RIRO traffic.

The response time, shown in figure 5.3, rises dramatically. It increases by about 10% after only 4 terminals have been

added. Doubling the number of terminals leads to a 60% increase in the response time.

Figures 5.4 to 5.6 show the results of increasing the number of BATCH jobs. Note that the MACRESPOND OCP utilization is not affected significantly. However, the response time is. The problem seems to lie in the exceptionally high utilization of SYSOO1 (almost 90% with 14 BATCH jobs).

The results of increasing the OCP speed is shown in figures 5.7 to 5.9. The OCP utilizations show an understandable drop. Note that the disc and disc controller utilizations will rise slightly because the jobs are not delayed for so long at the OCP. The response time drops by 10% when the OCP speed is doubled. It is doubtful that this would be noticed by the terminal user.

5.2.2 Configuration with disc traffic equalized

The graphs in figures 5.10 to 5.18 correspond to those in figures 5.1 to 5.9. The characteristics are the same as discussed in section 5.2.1. The following points should be noted:

- the removal of the disc bottleneck allows more work to be performed - i.e. the OCP utilizations will increase;

- the utilization of SYSOO1 drops dramatically (note that all discs will have the same utilization);

- the utilization of DC40 rises, while that of DC10 drops. This is due to the fact that traffic that previously went through DC10 is now routed through DC40. DC40 also has six discs attached to it, compared with the five on DC10. Finally, DC40 also has the RIRO traffic. A pure equalization of the disc file traffic is therefore not

the optimal solution. One can (and should) perform further experiments. This is the beauty of the model. One can make numerous predictions without ever having to interfere with the live system;

- the response times are all consistently lower than the corresponding examples in section 5.2.1.
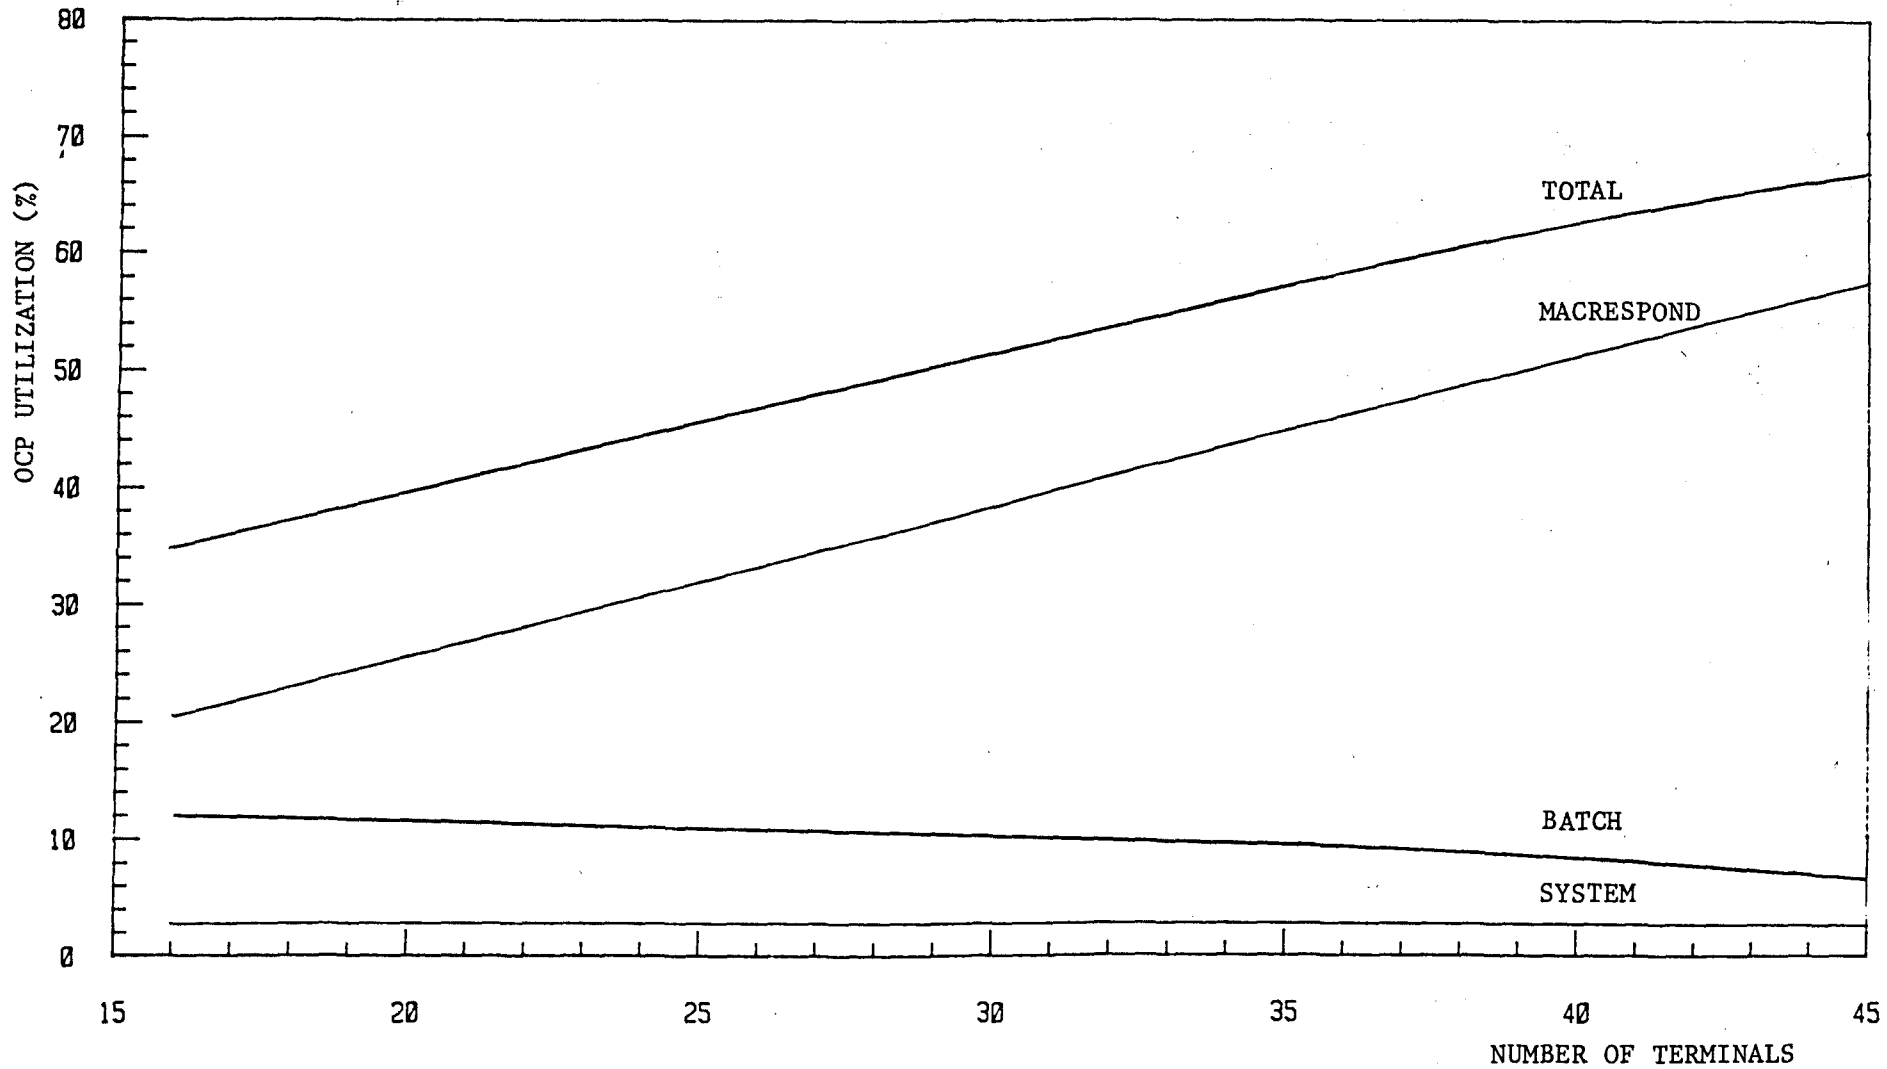
Figure 5.1   JUNO41 configuration. Number of terminals vs OCP utilization.
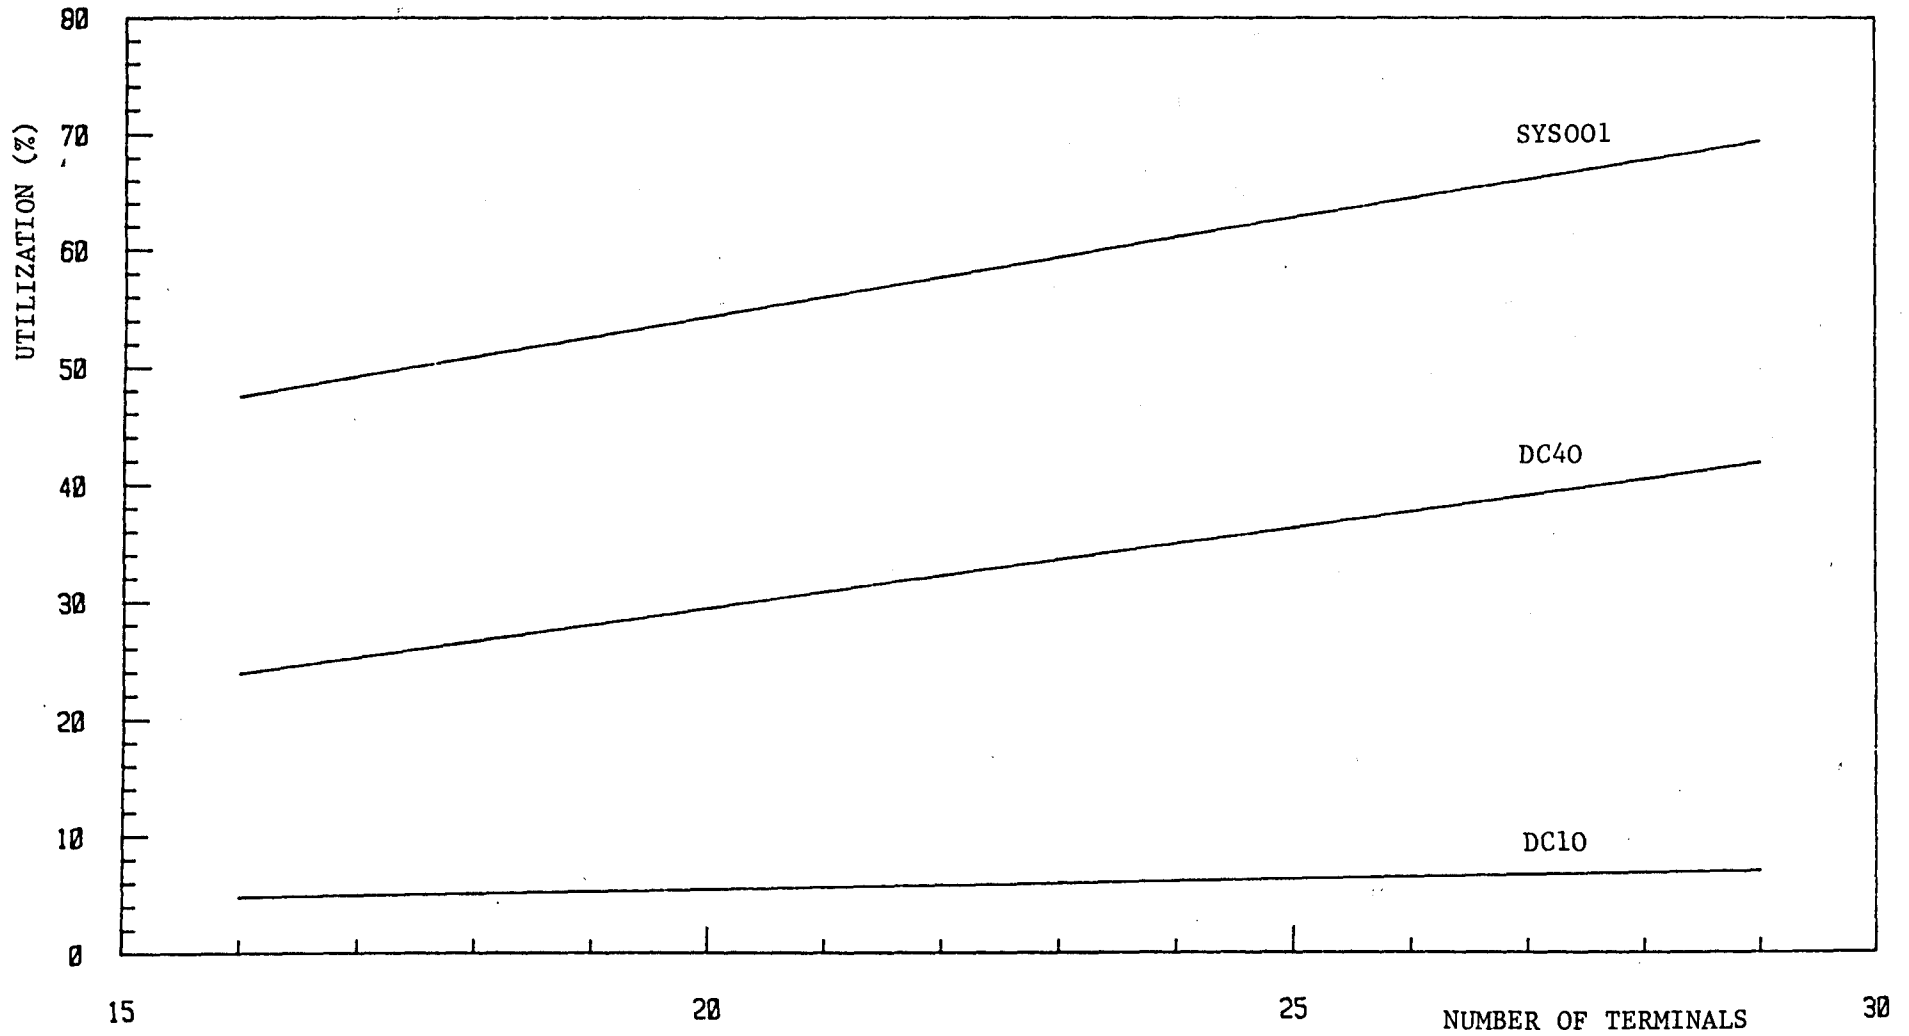
UTILIZATION (%)

80

70

60

50

40

30

20

10

0

SYS001

DC40

DC10

15

20

25

30

NUMBER OF TERMINALS

Figure 5.2   JUN041 configuration. Number of terminals vs SYS001 and disc controller utilizations.
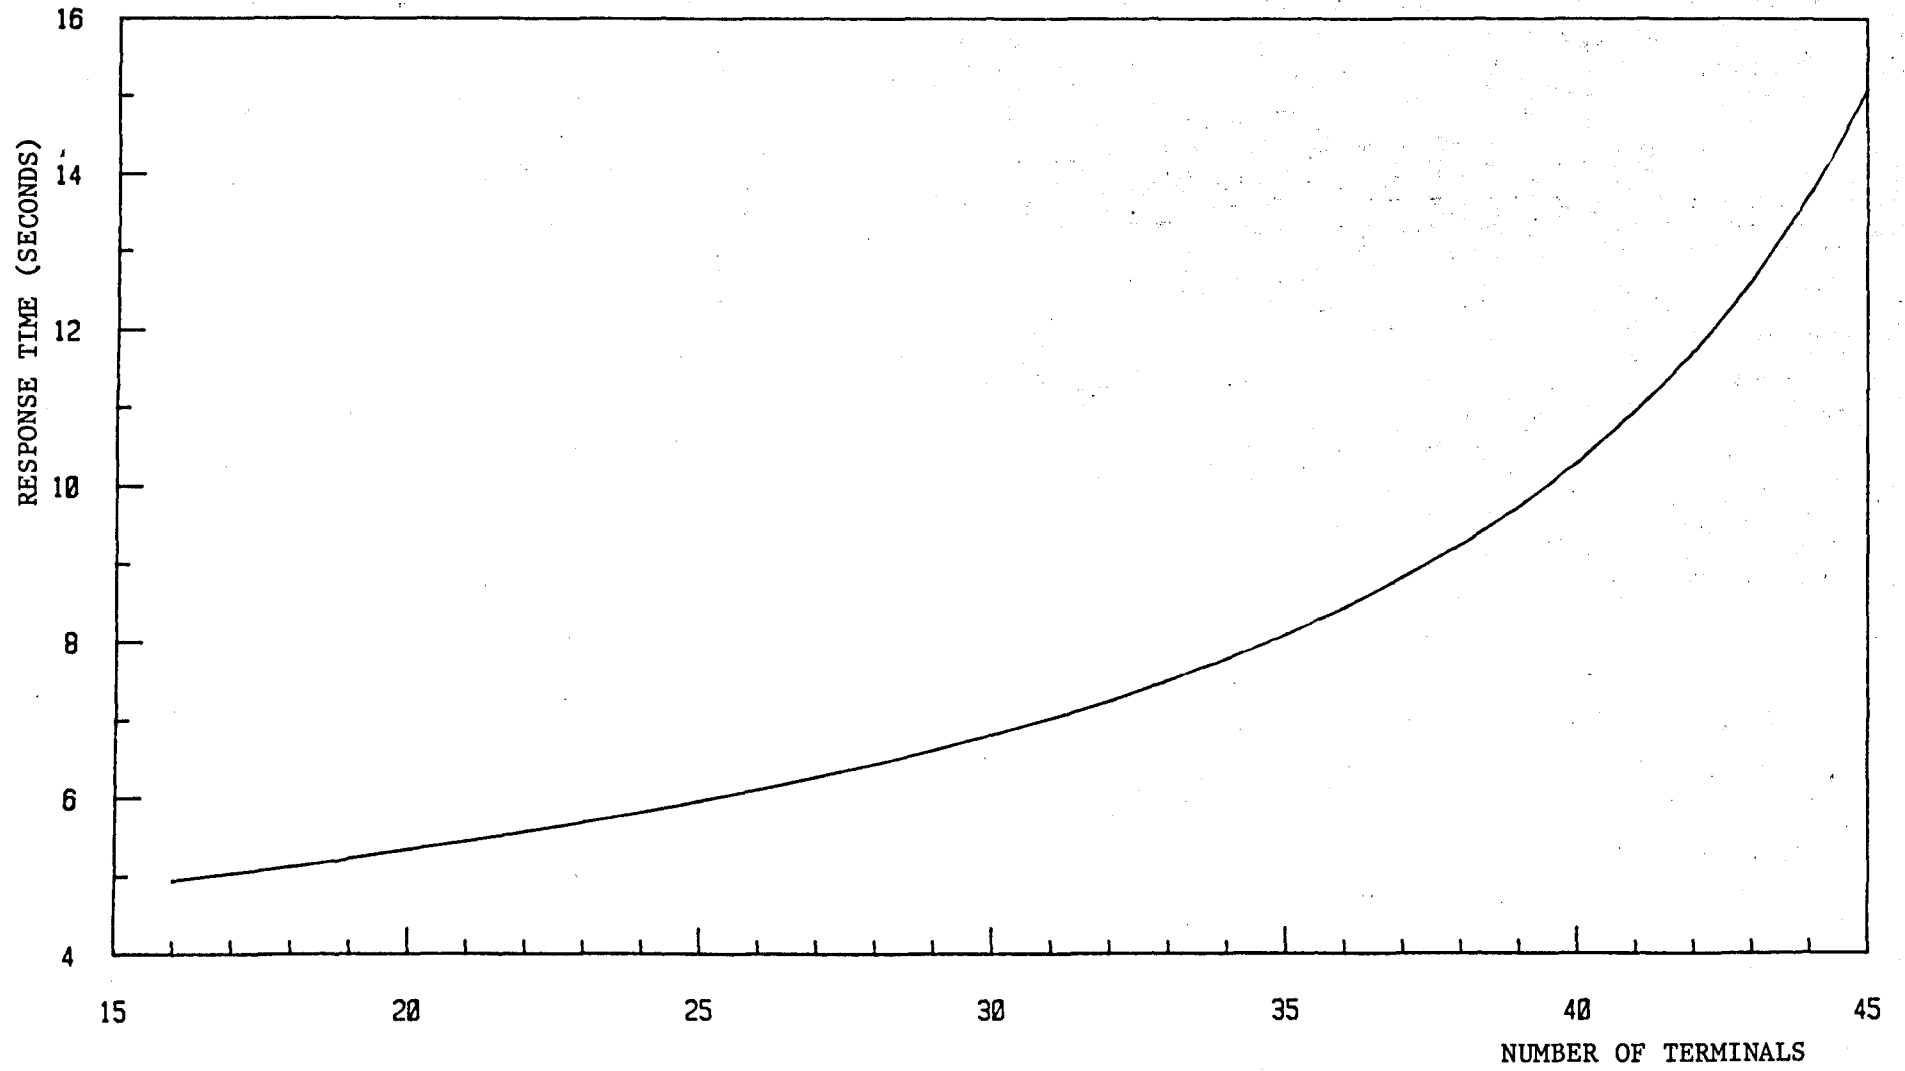
Figure 5.3   JUNO41 configuration. Number of terminals vs response time.
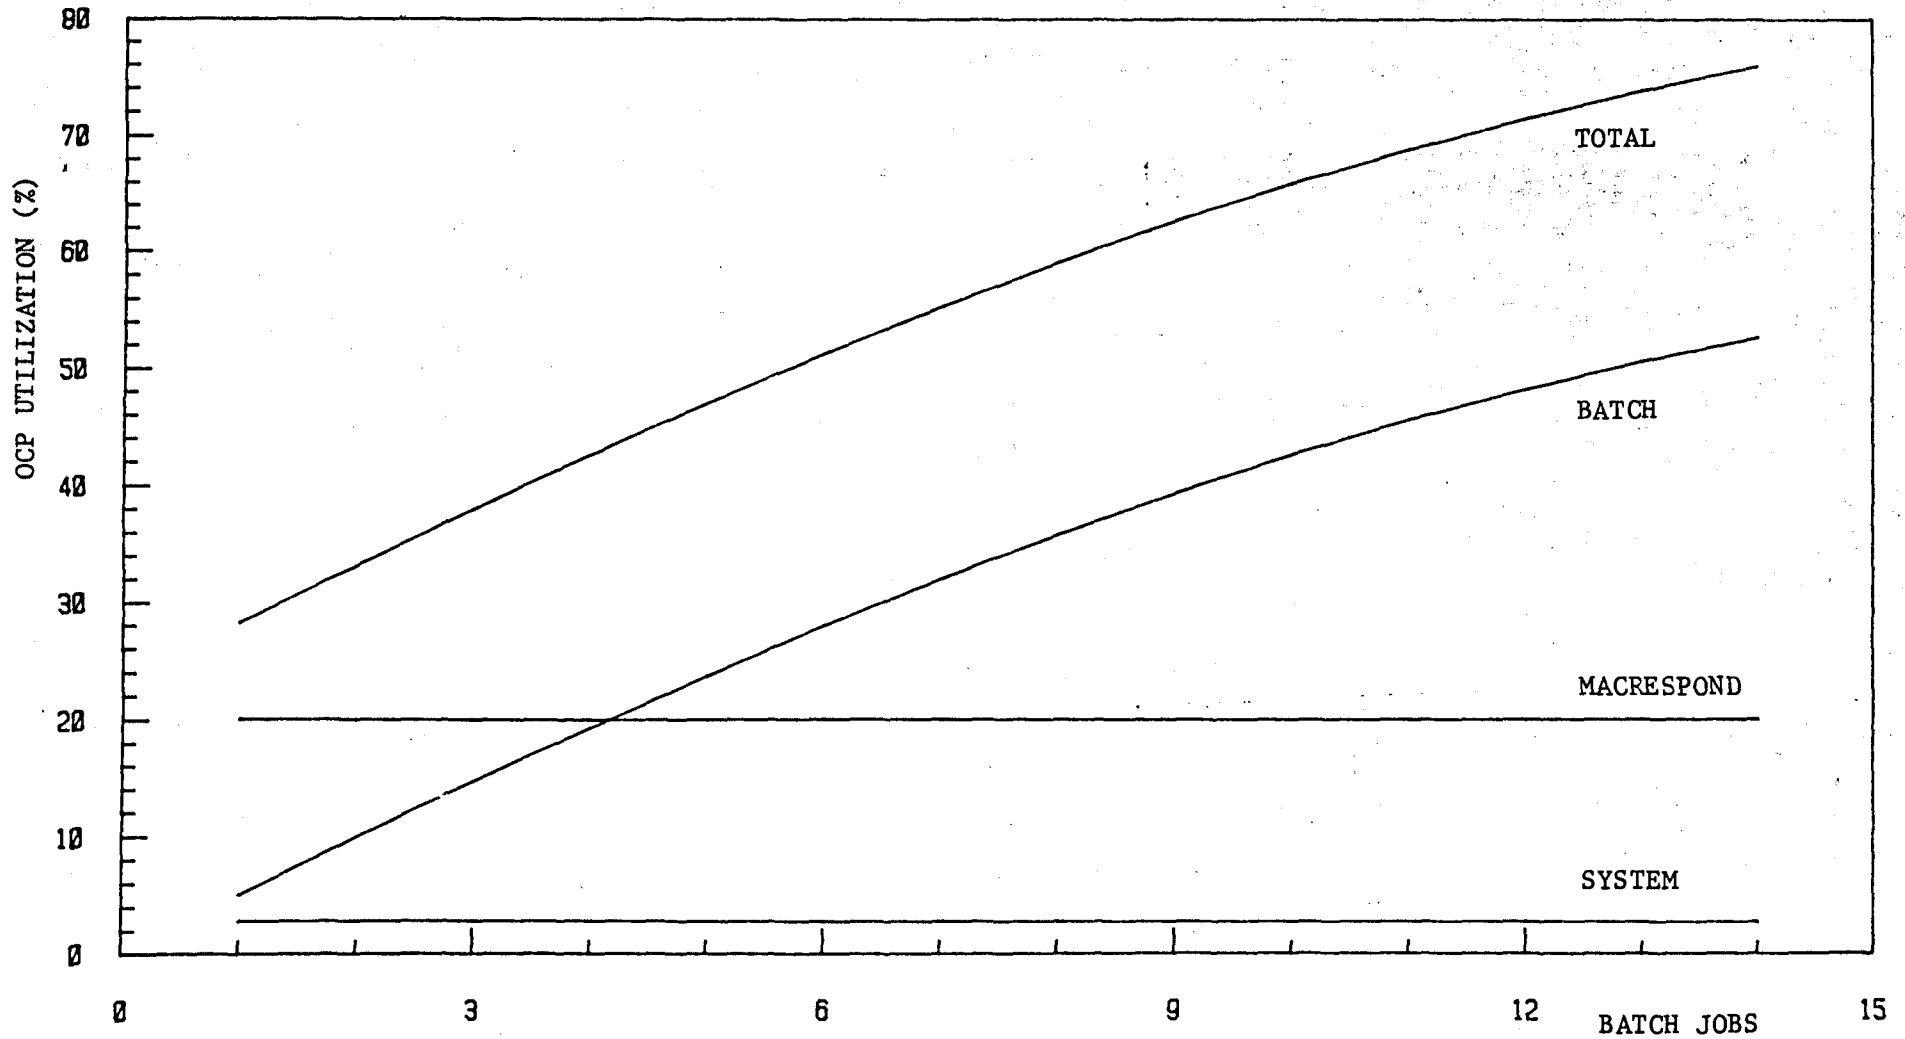
Figure 5.4   JUN041 configuration. BATCH jobs vs OCP utilization.
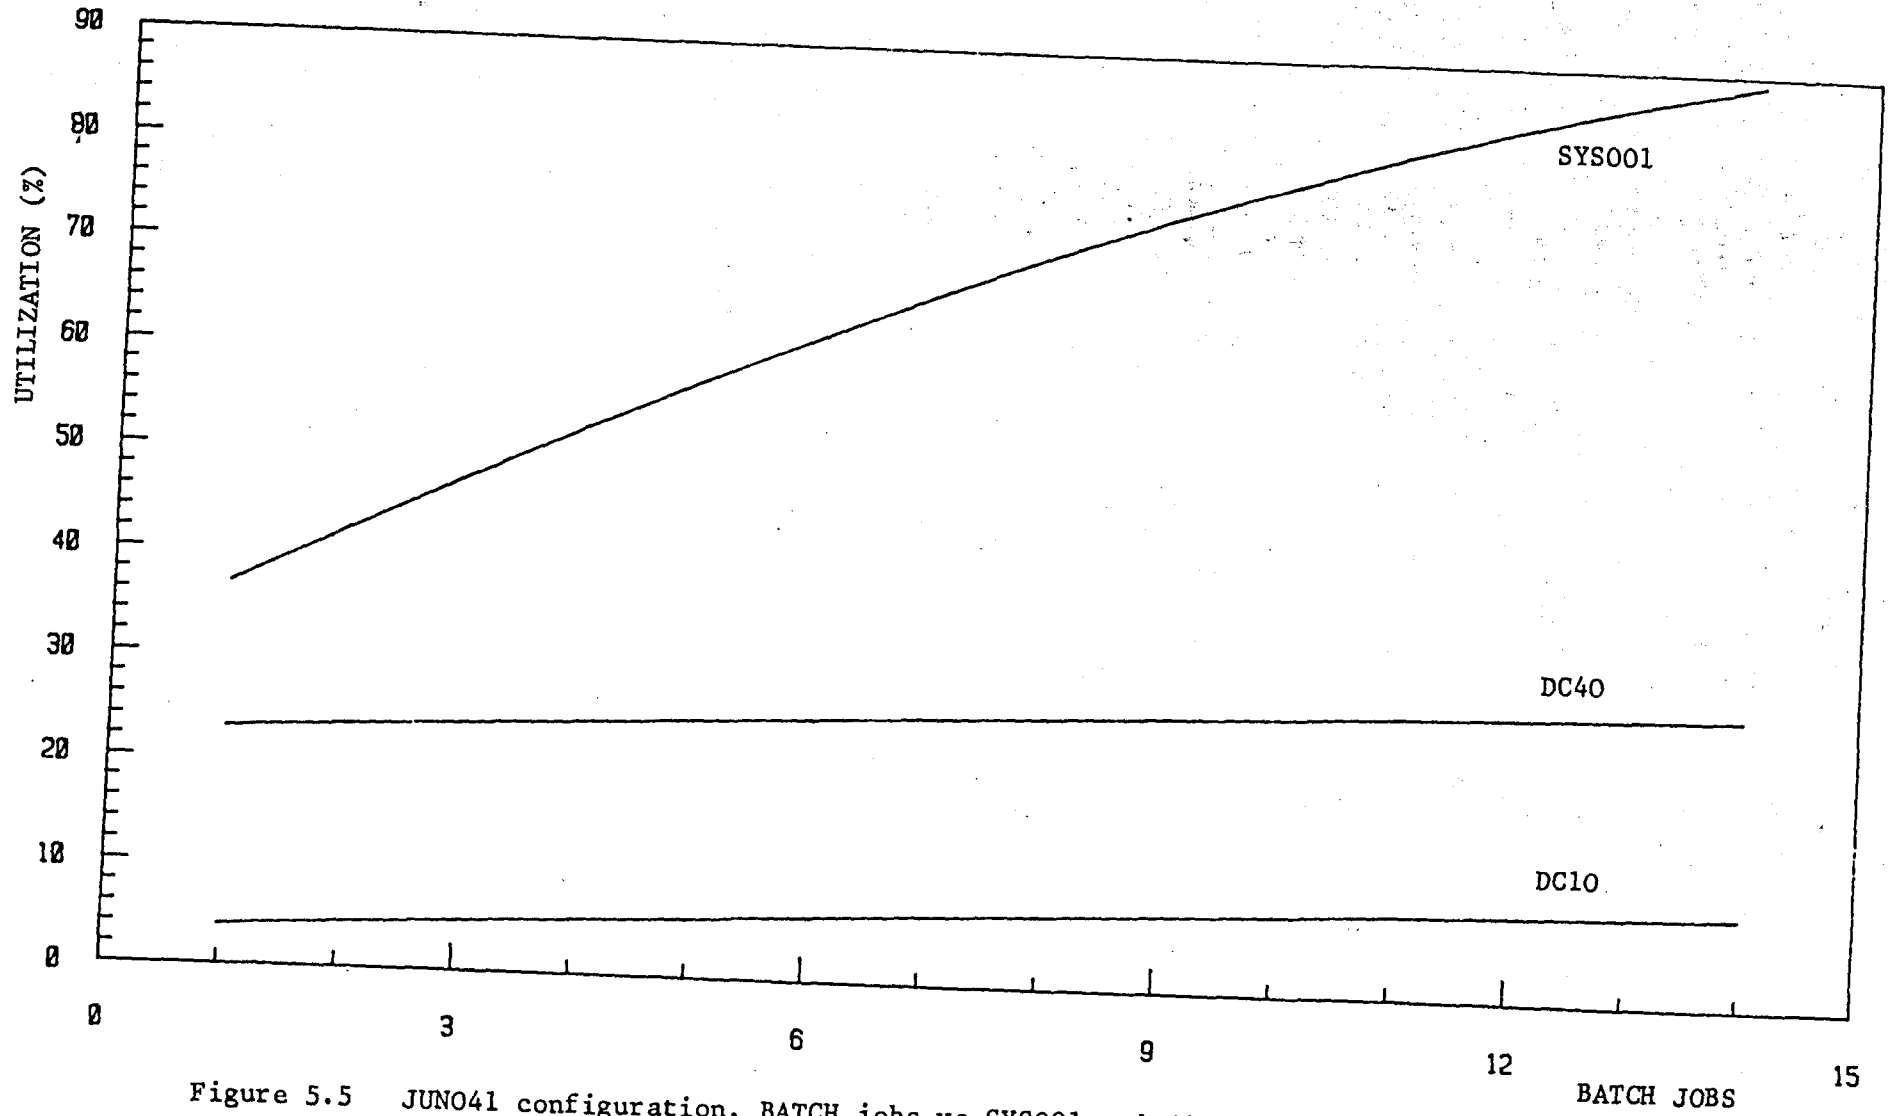
Figure 5.5   JUN041 configuration. BATCH jobs vs SYS001 and disc controller utilizations.
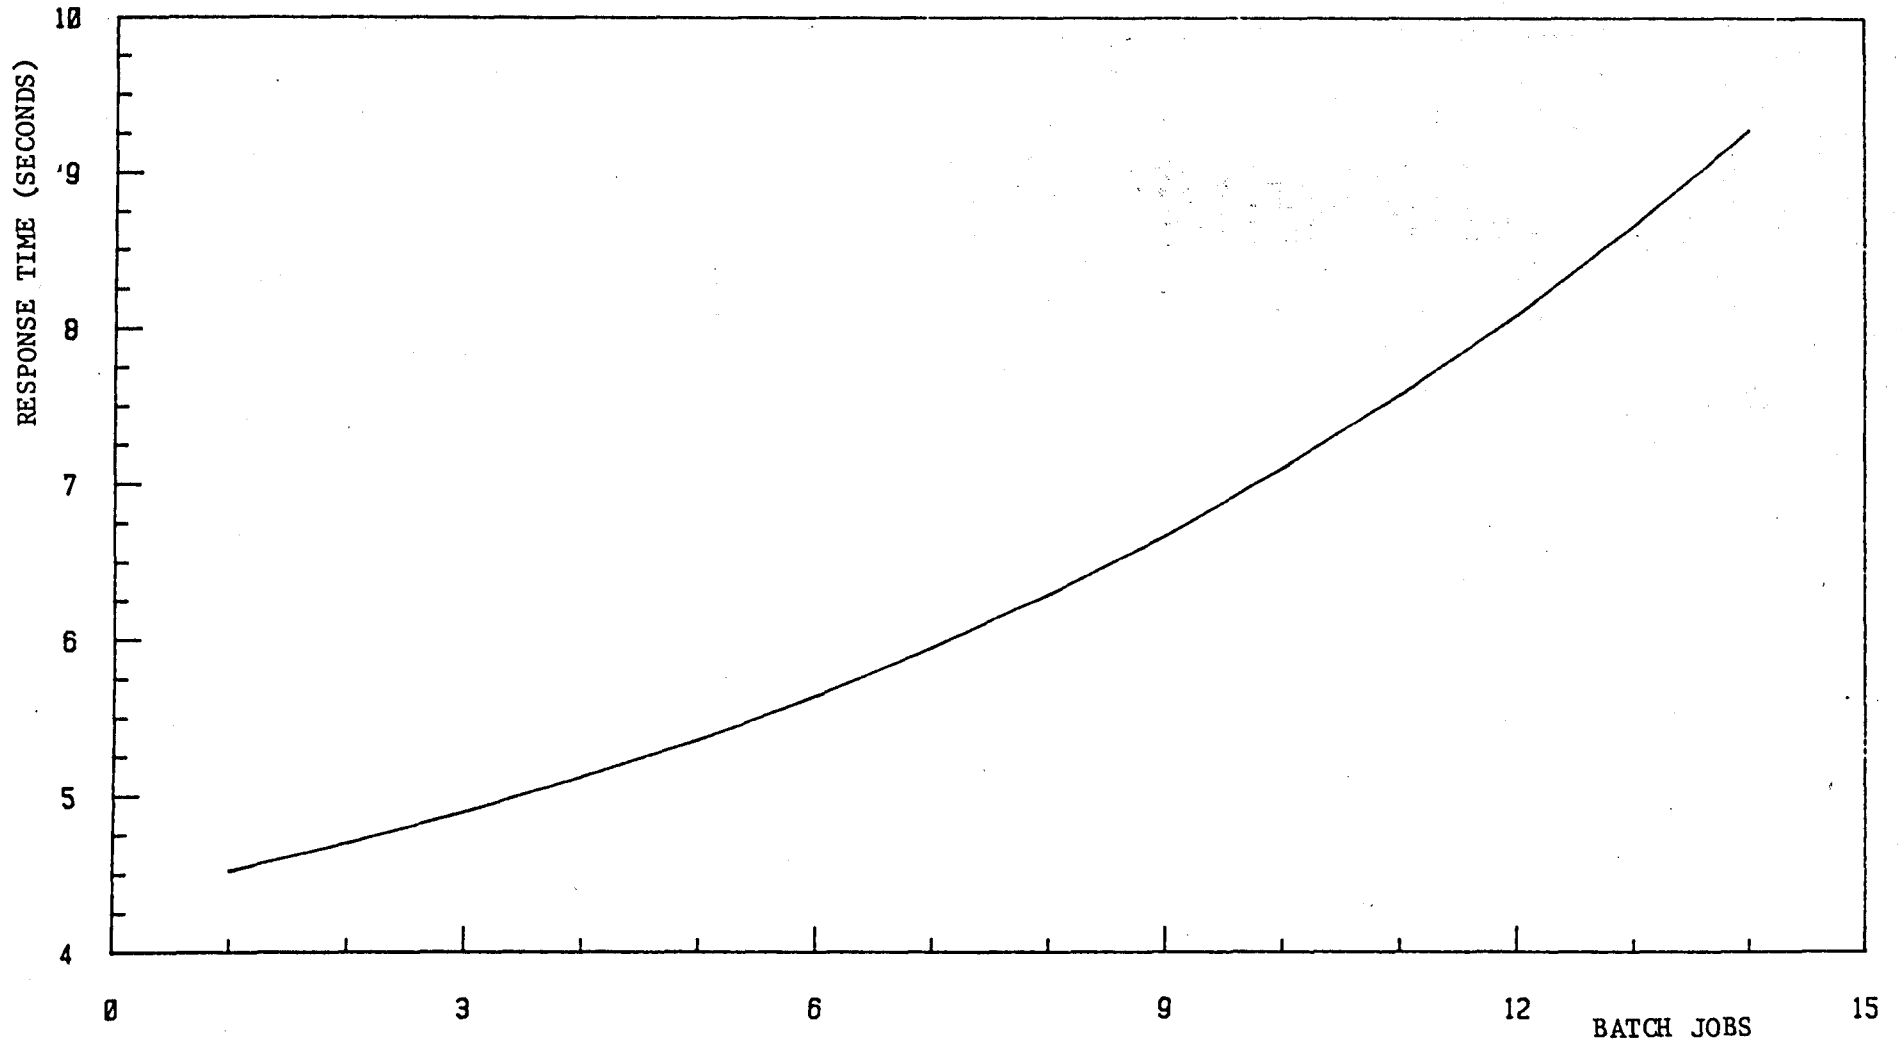
Figure 5.6   JUNO41 configuration. BATCH jobs vs response time.
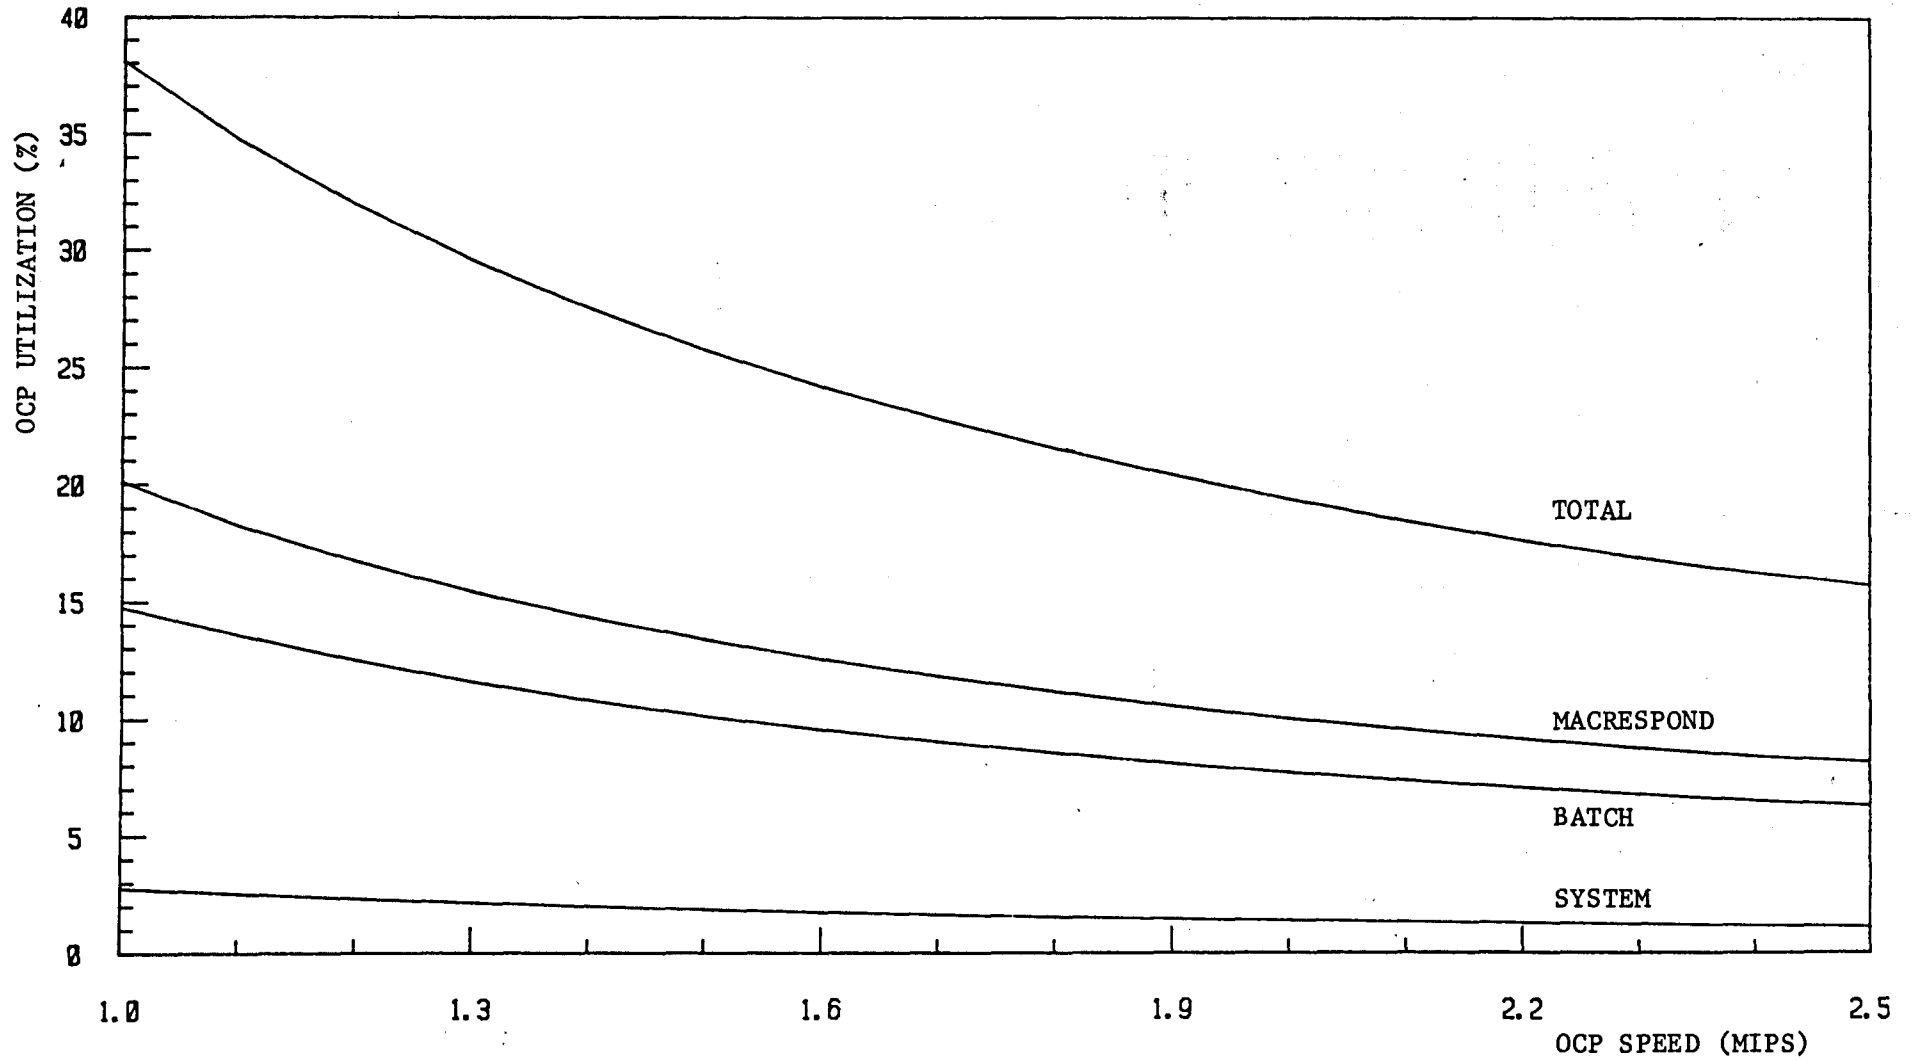
Figure 5.7   JUN041 configuration. OCP speed vs OCP utilization.
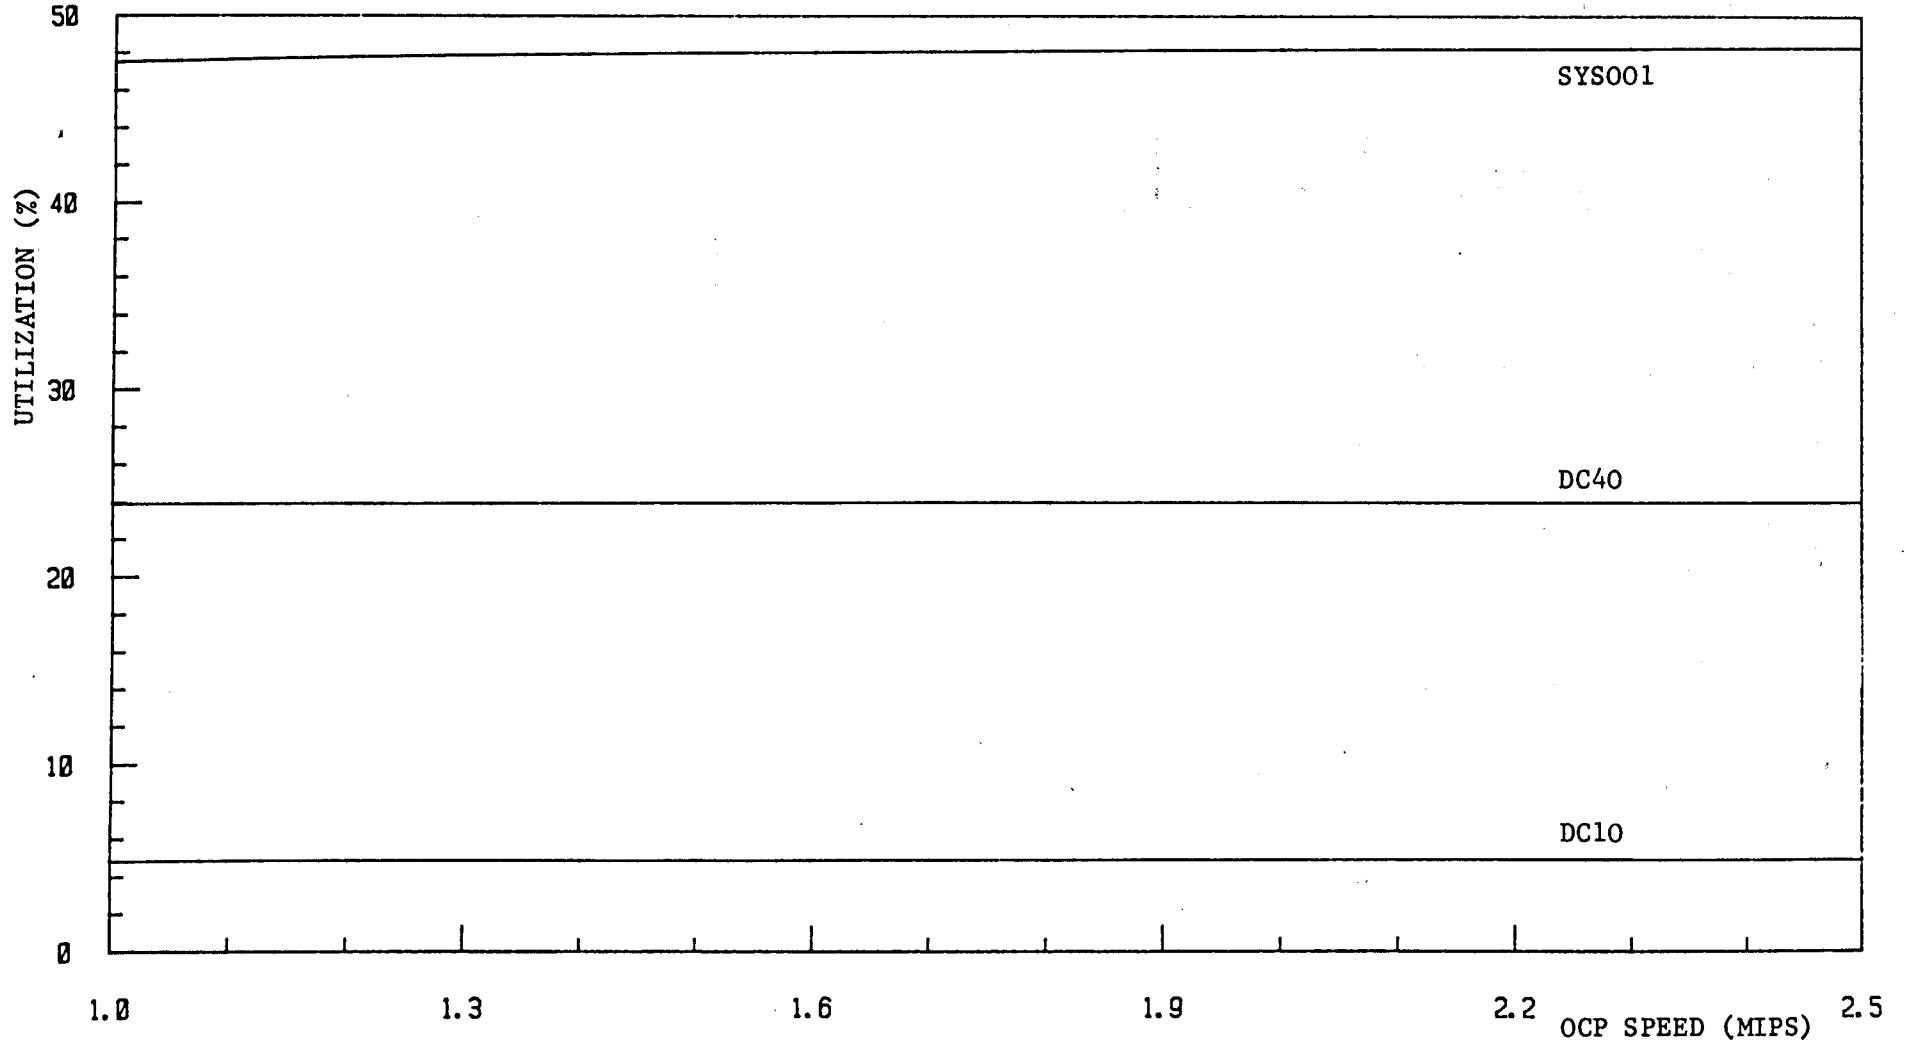
Figure 5.8   JUN041 configuration. OCP speed vs SYS001 and disc controller utilizations.
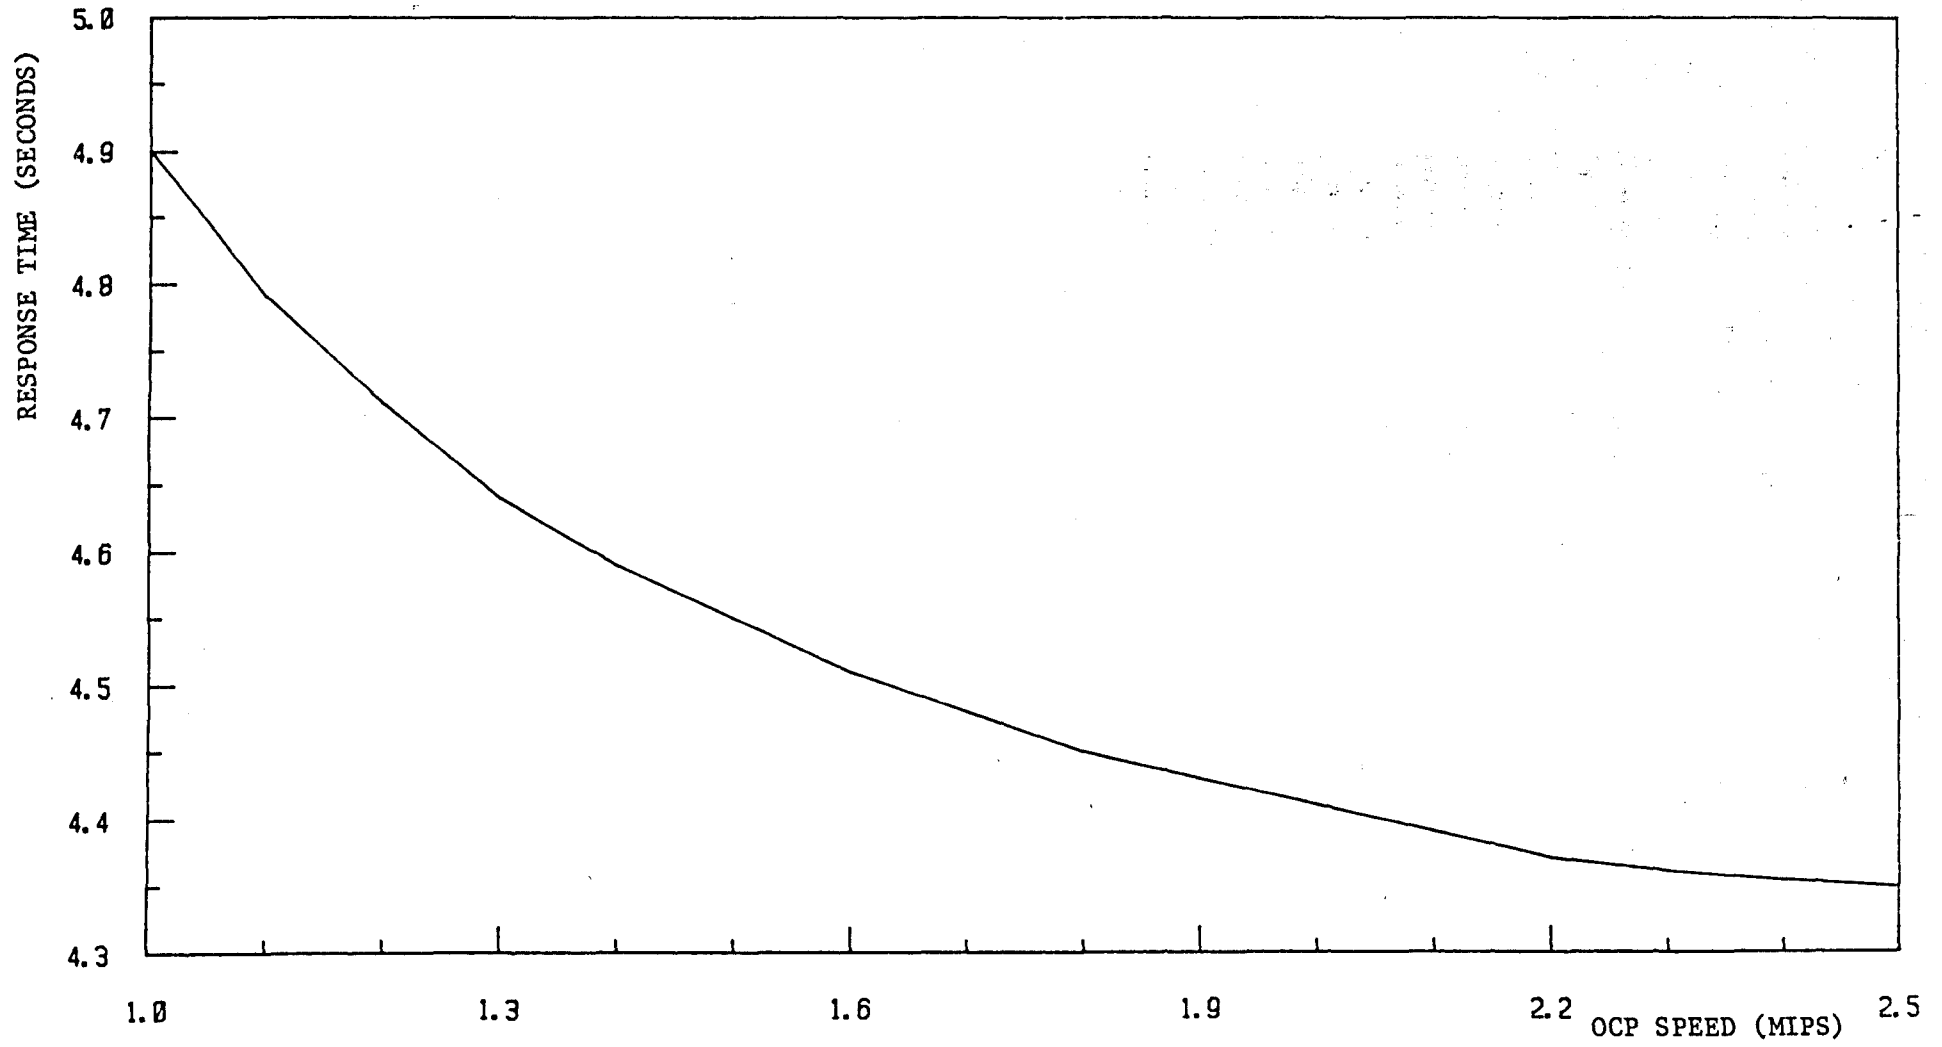
Figure 5.9   JUNO41 configuration. OCP speed vs response time.
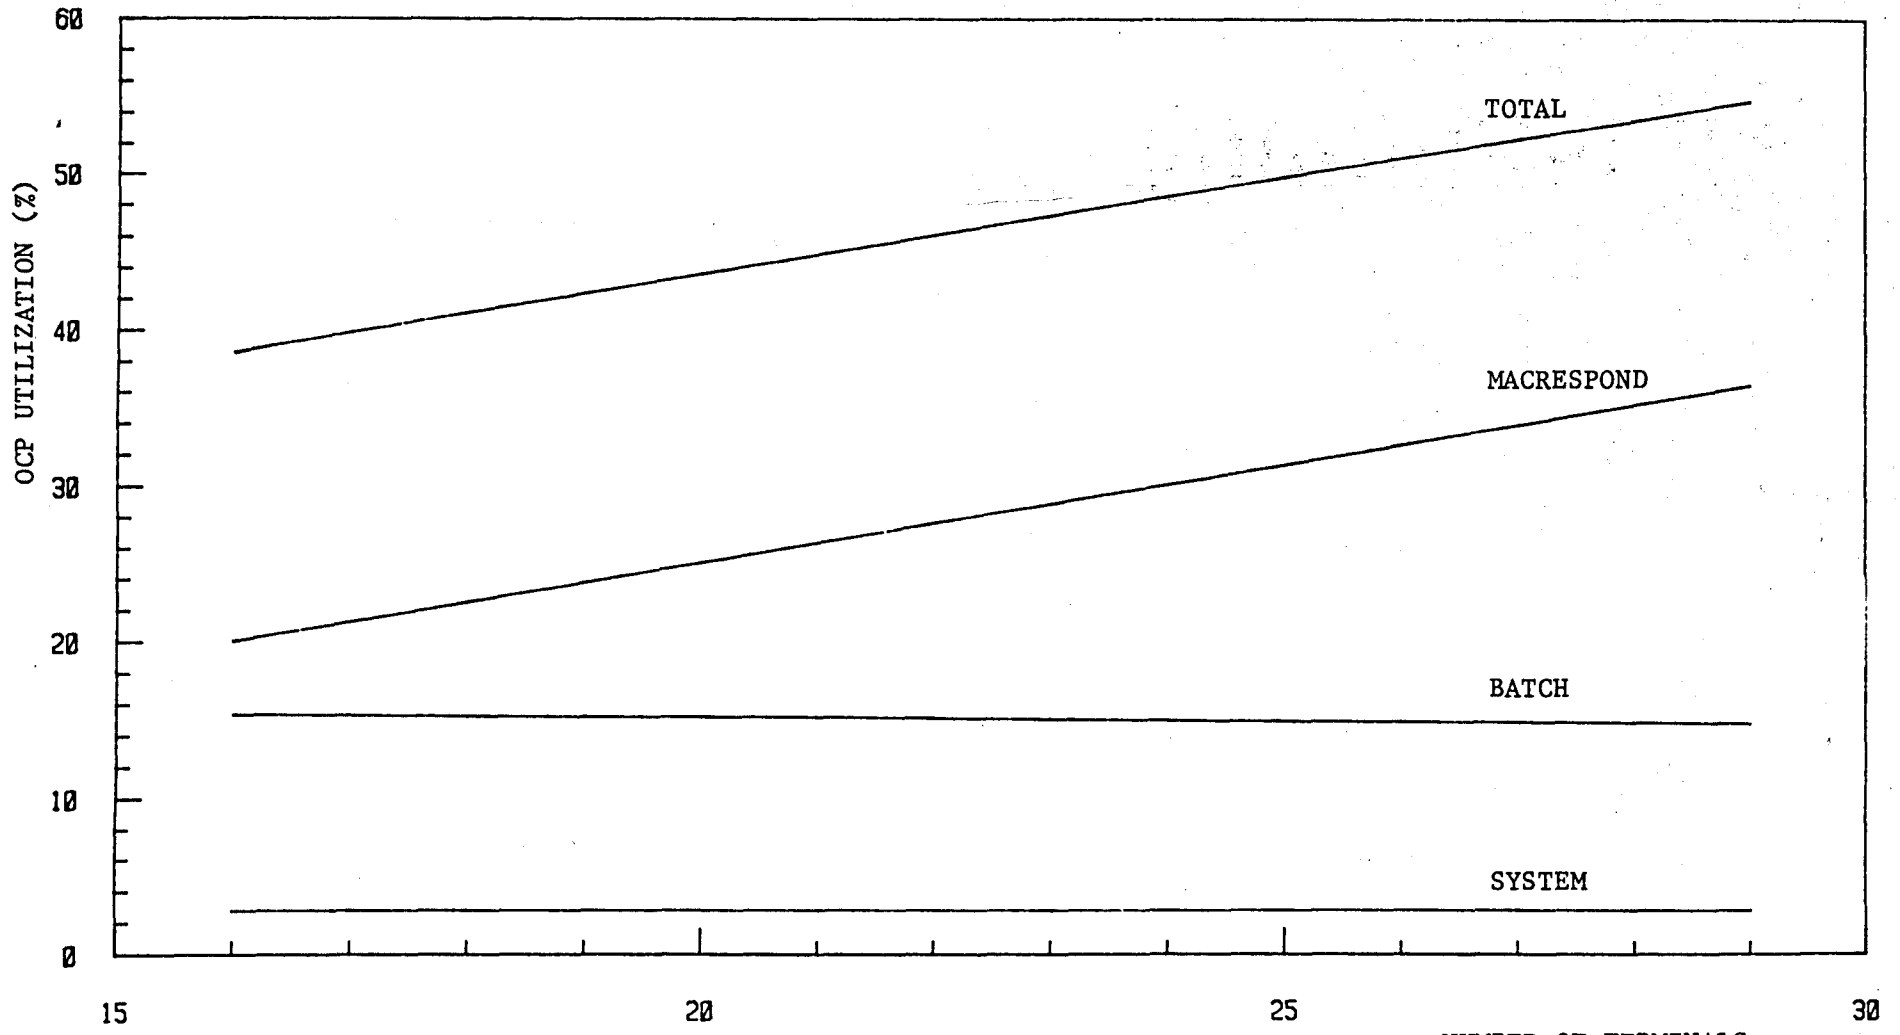
Figure 5.10   Disc traffic equalized. Number of terminals vs OCP utilization.

110
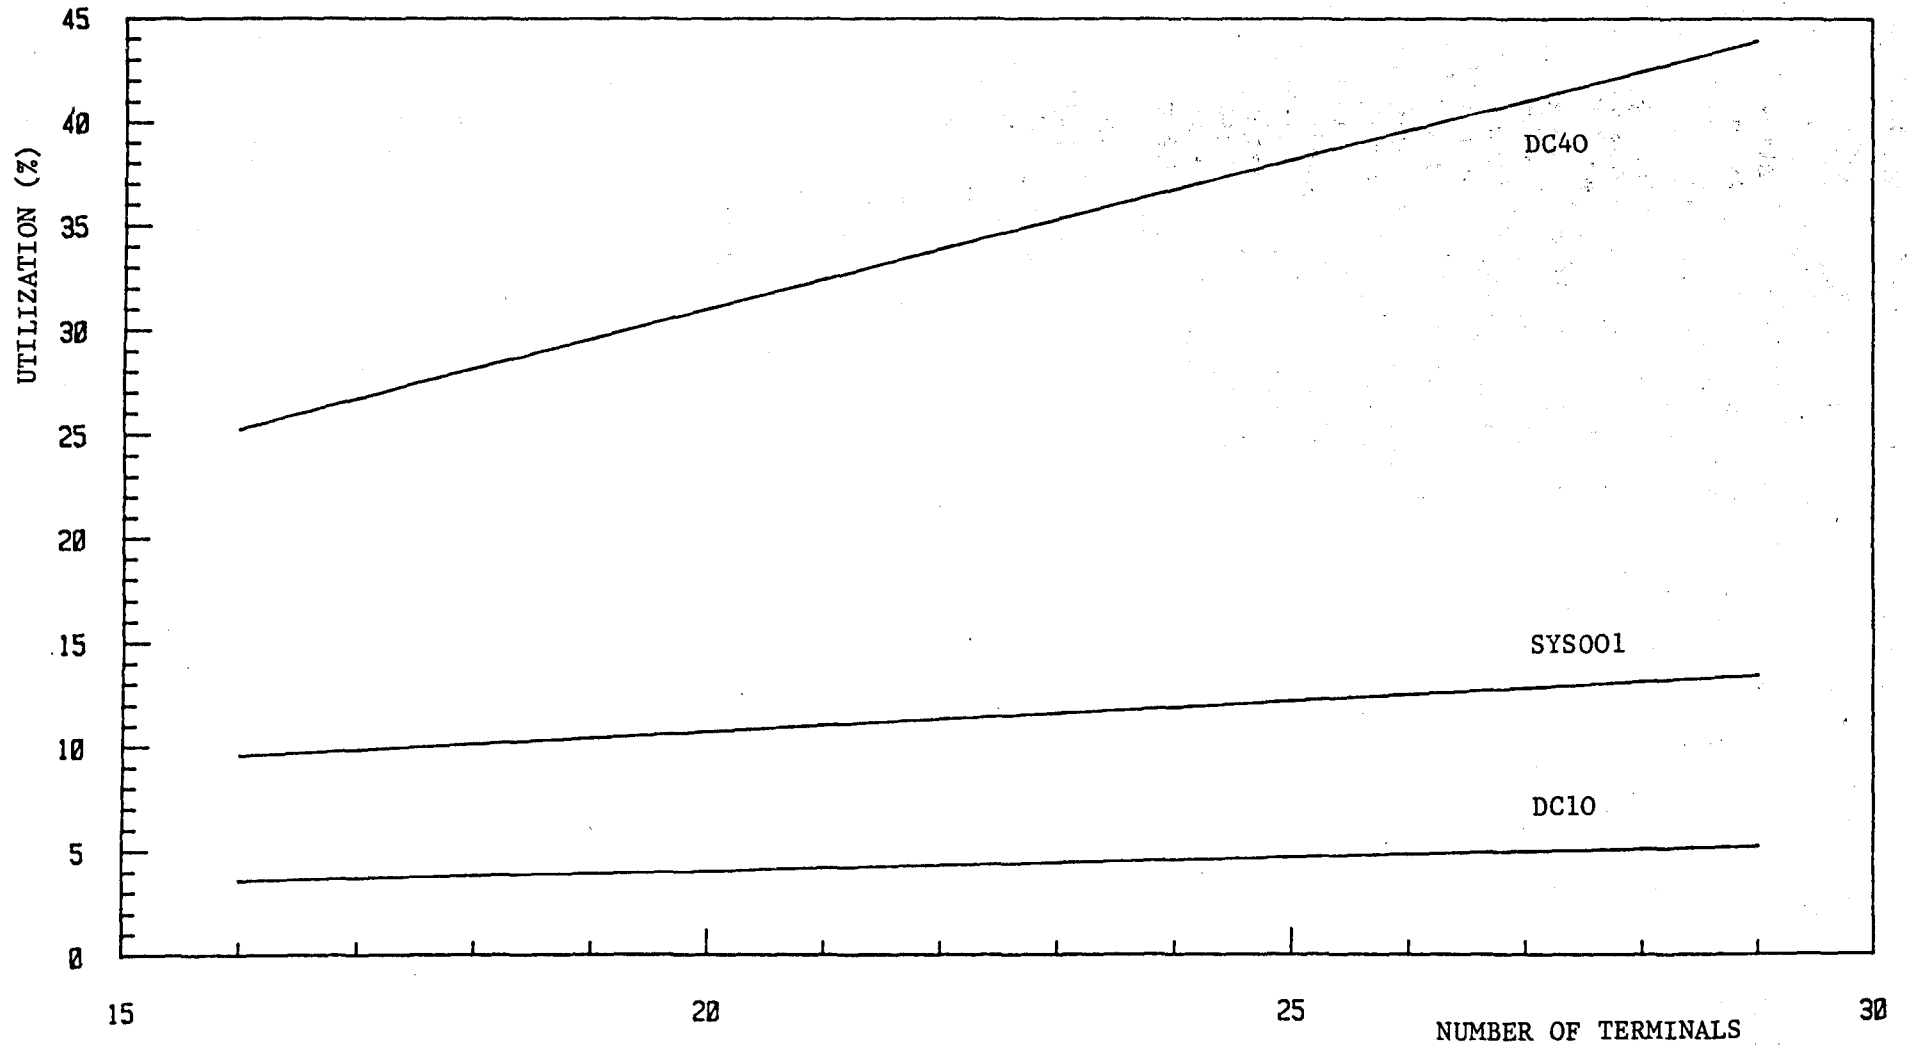
Figure 5.11   Disc traffic equalized. Number of terminals vs SYS001 and disc controller utilizations.
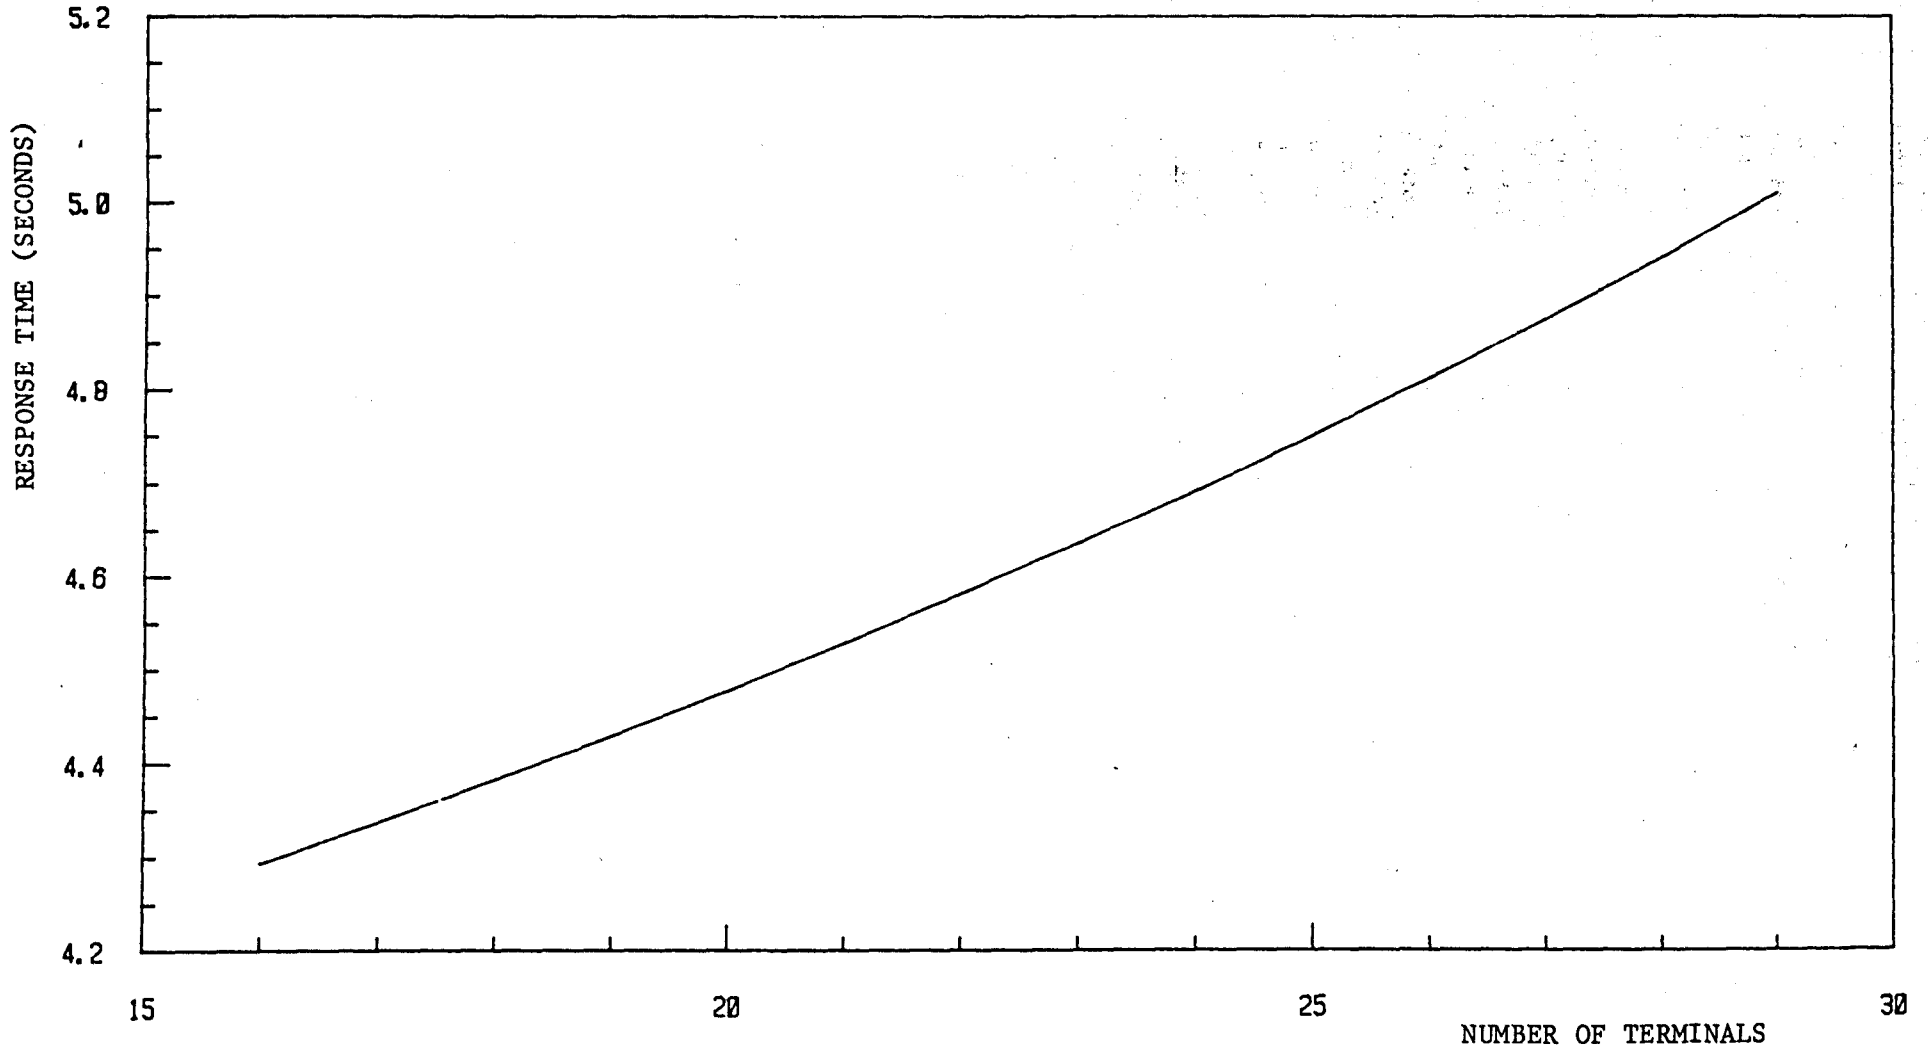
Figure 5.12   Disc traffic equalized. Number of terminals vs response time.
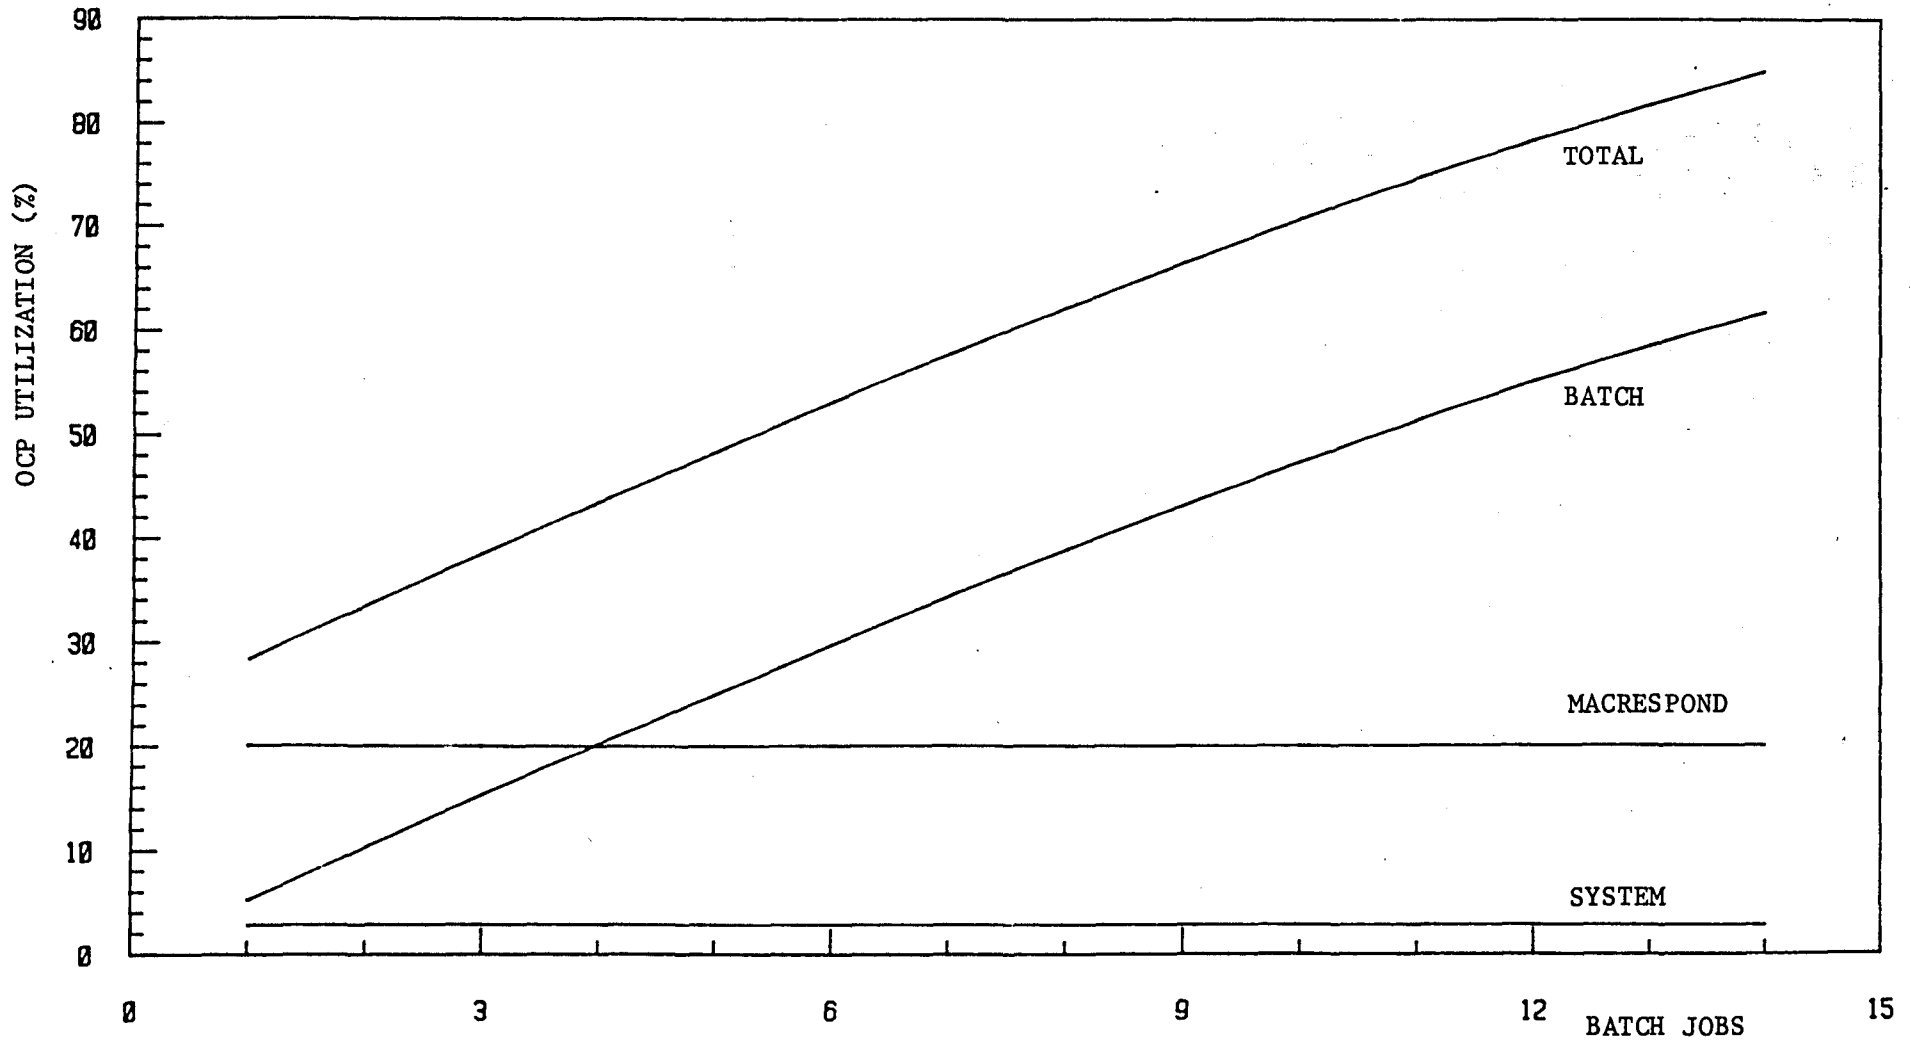
Figure 5.13   Disc traffic equalized. BATCH jobs vs OCP utilization.

113
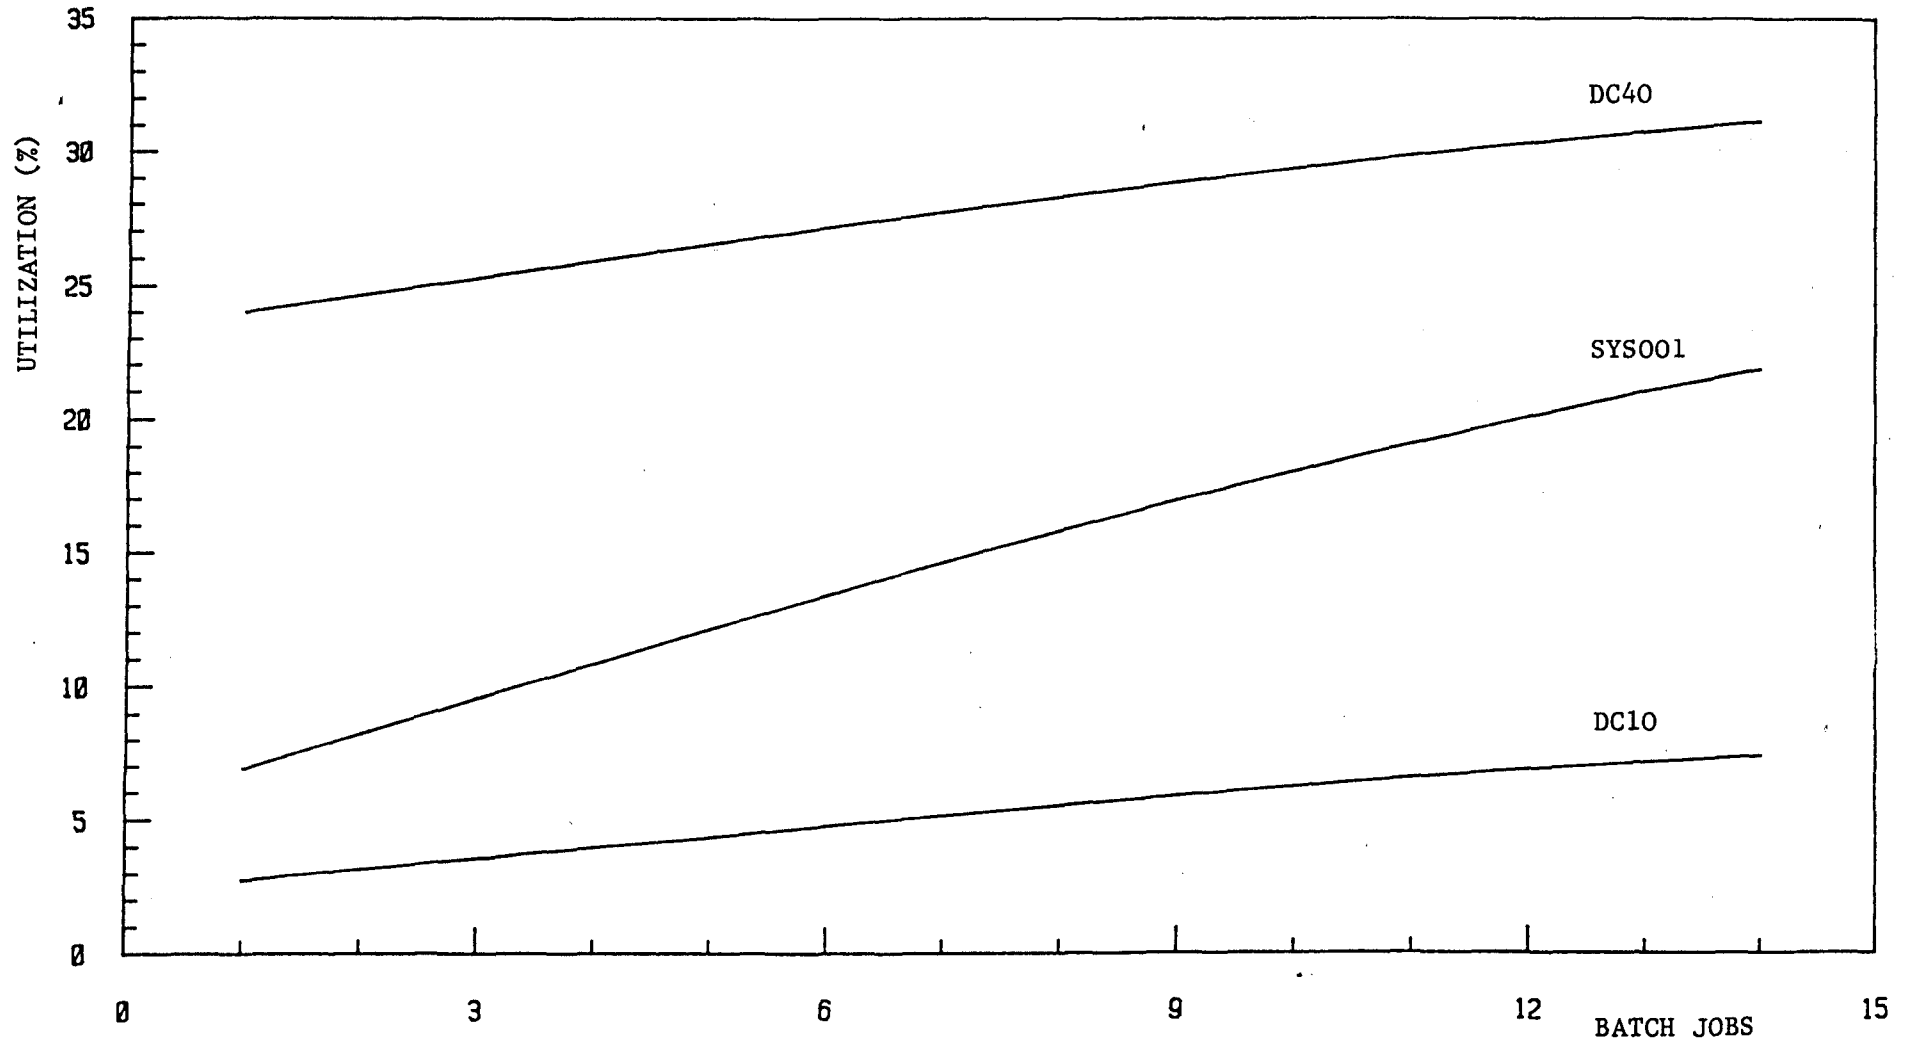
Figure 5.14   Disc traffic equalized. BATCH jobs vs SYS001 and disc controller utilizations.
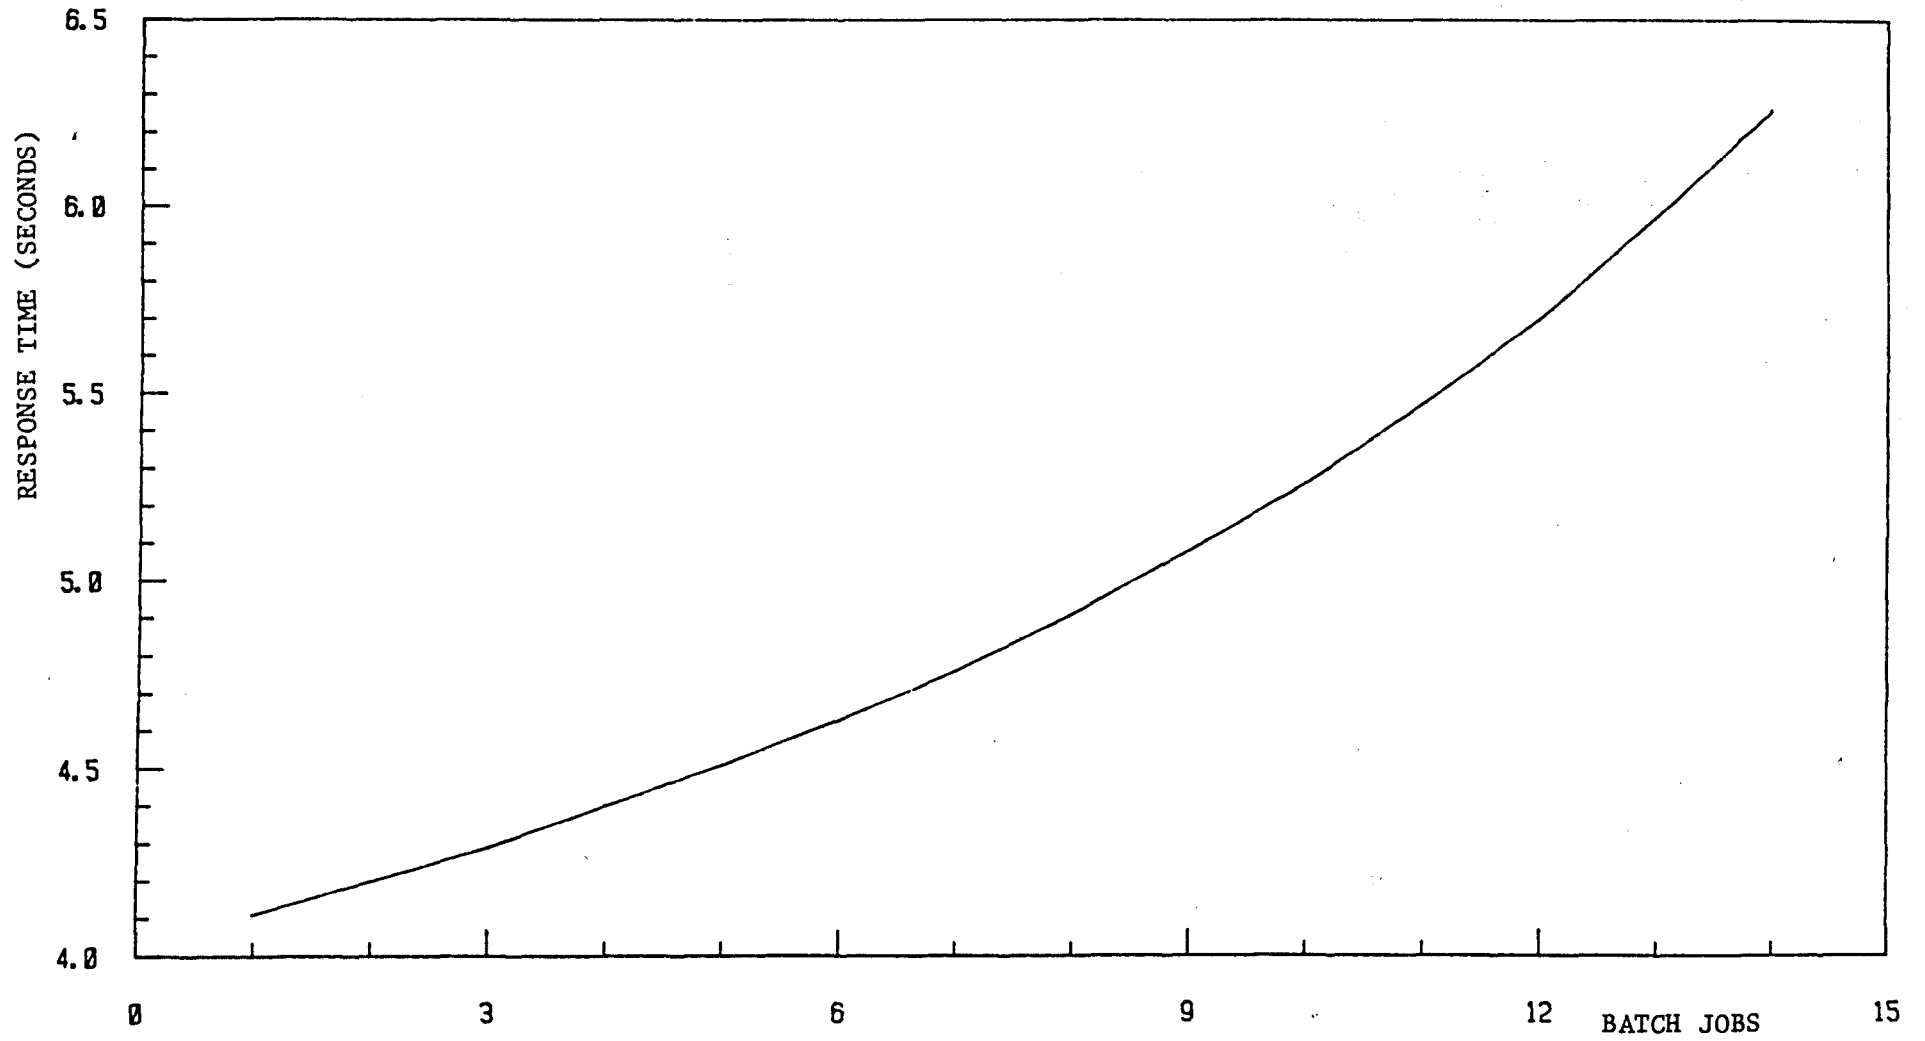
Figure 5.15   Disc traffic equalized. BATCH jobs vs response time.
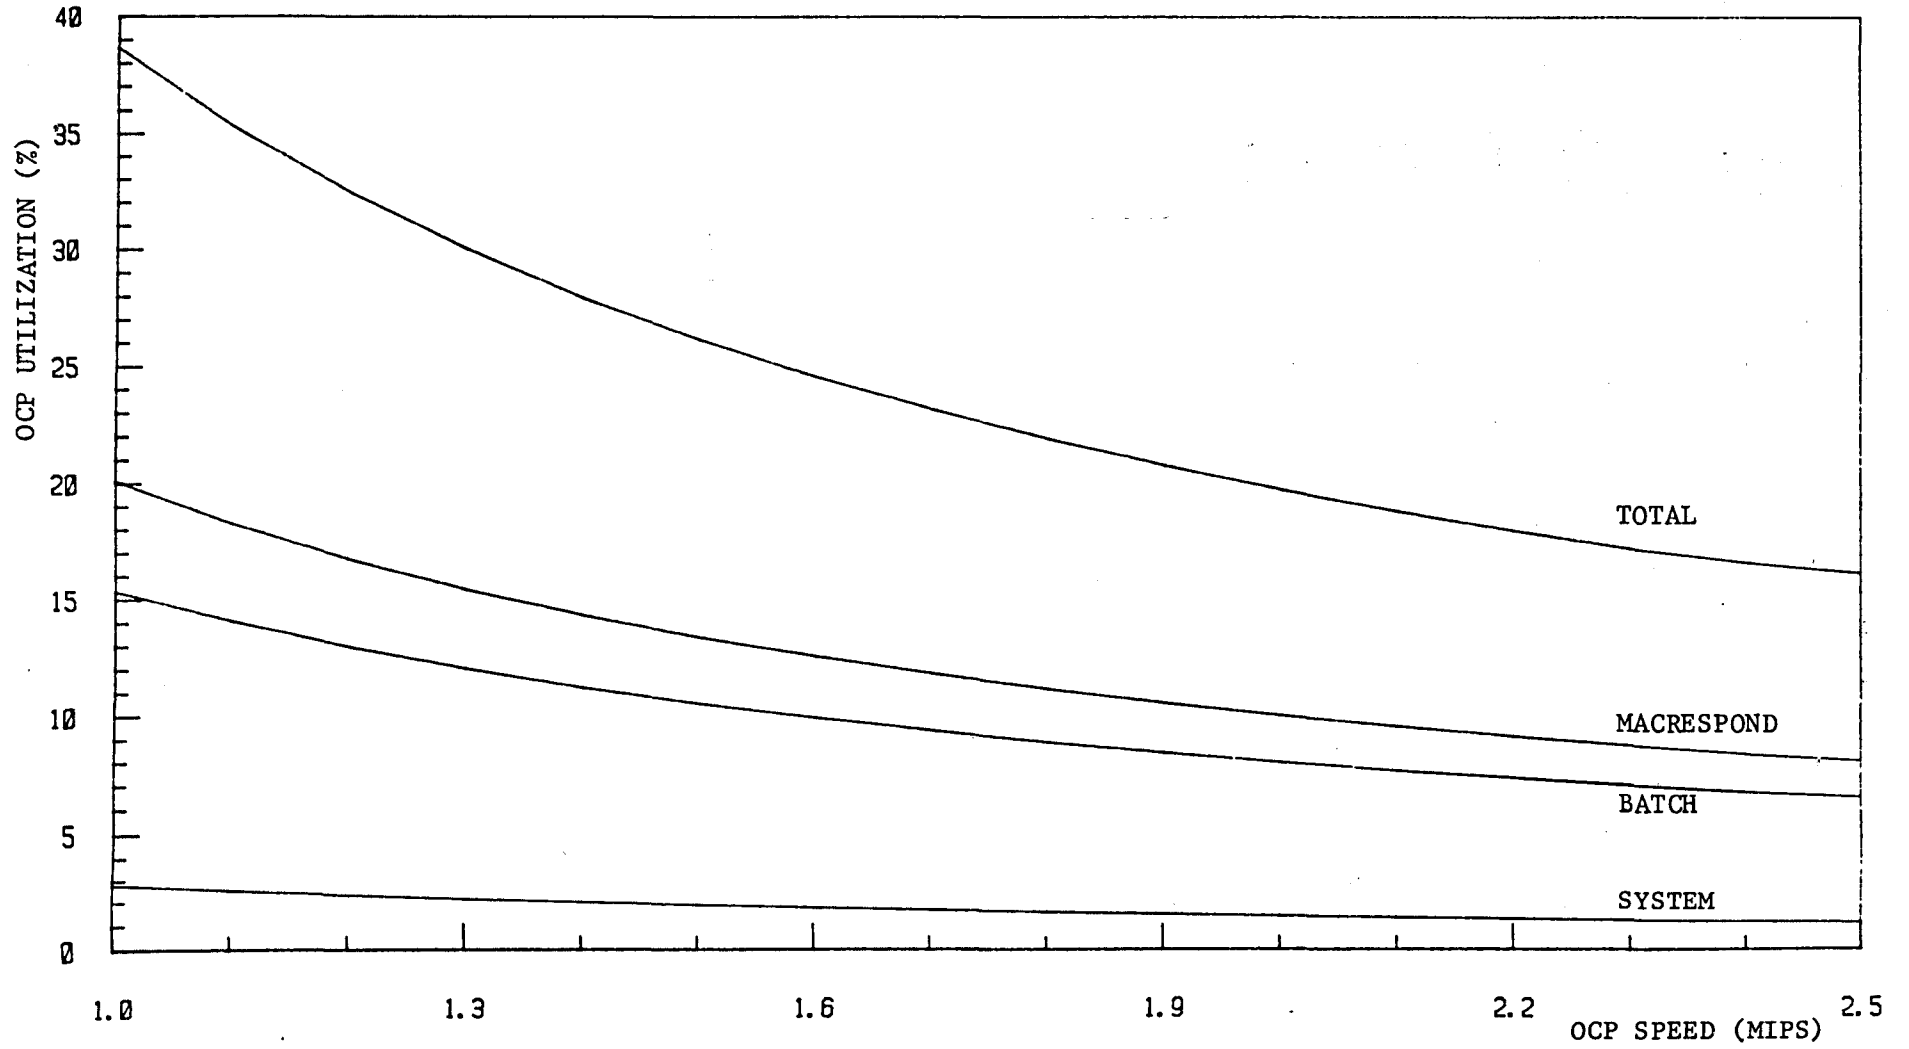
Figure 5.16   Disc traffic equalized. OCP speed vs OCP utilization.
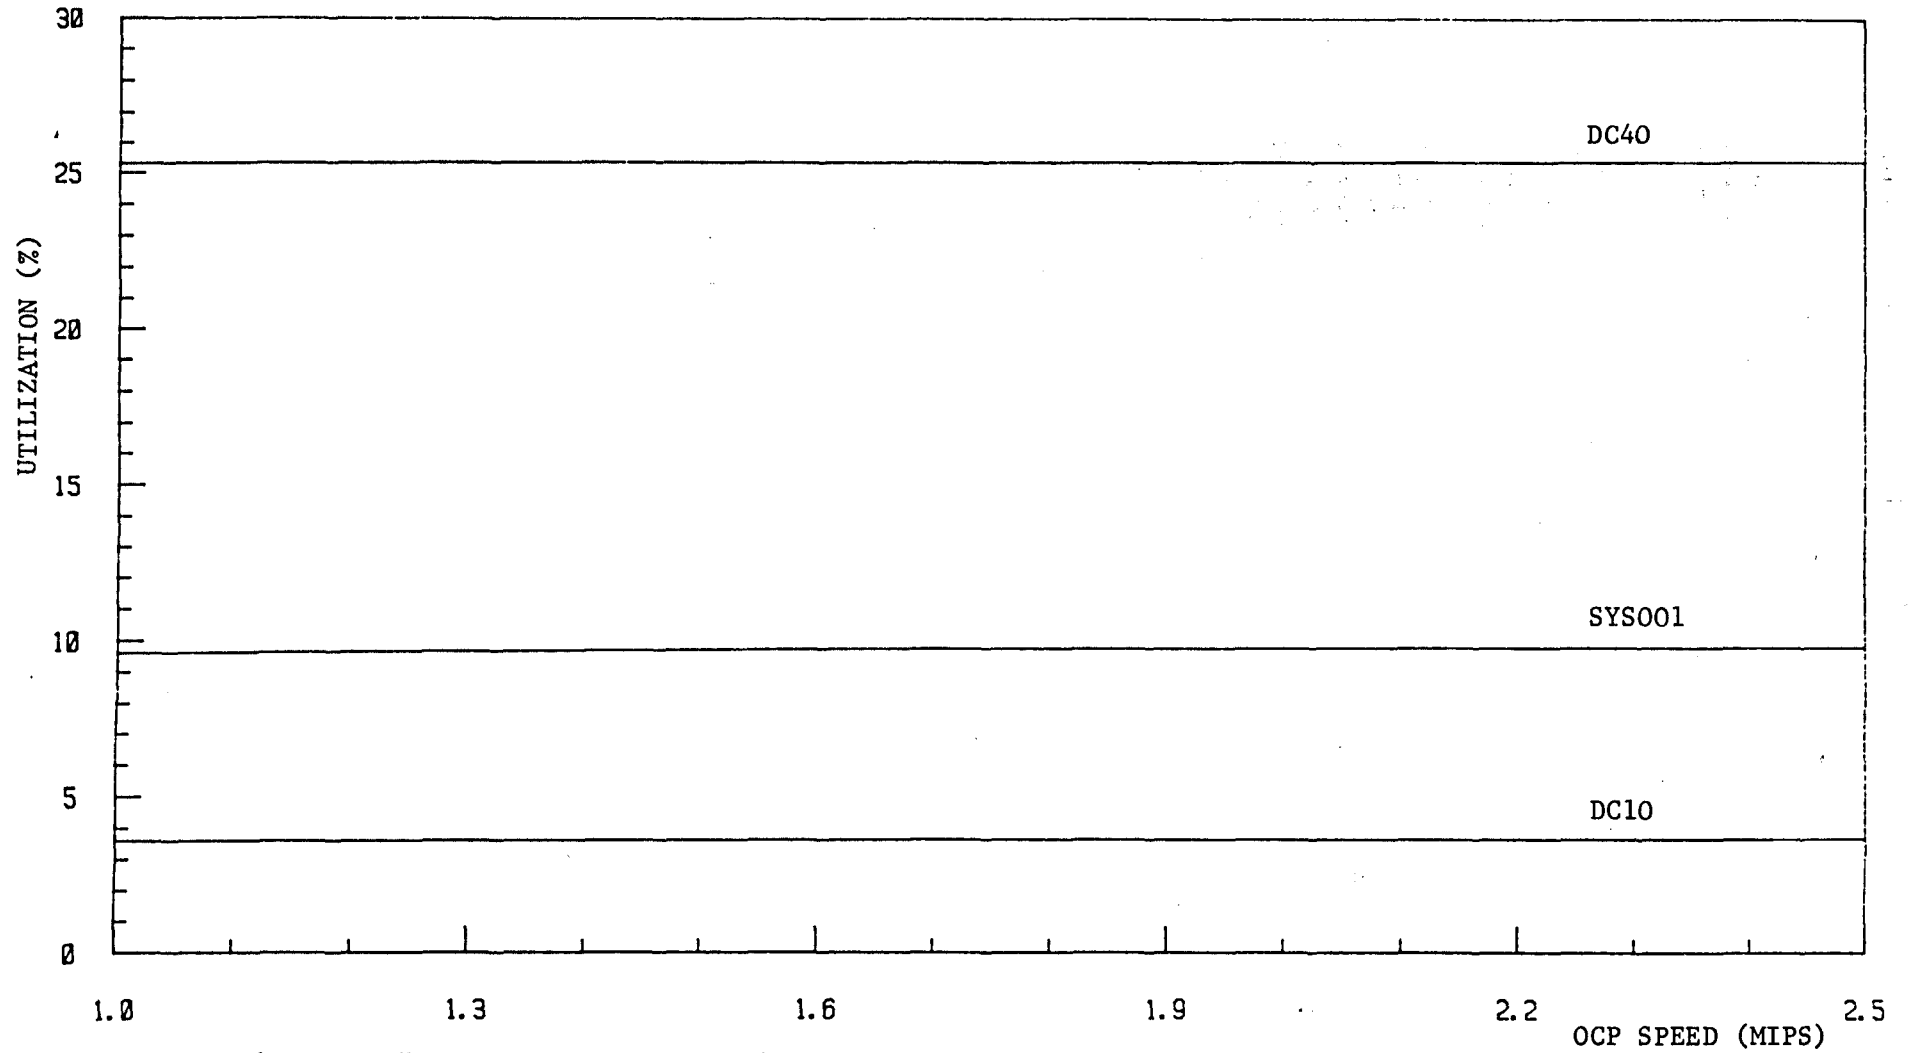
Figure 5.17   Disc traffic equalized. OCP speed vs SYS001 and disc controller utilization.
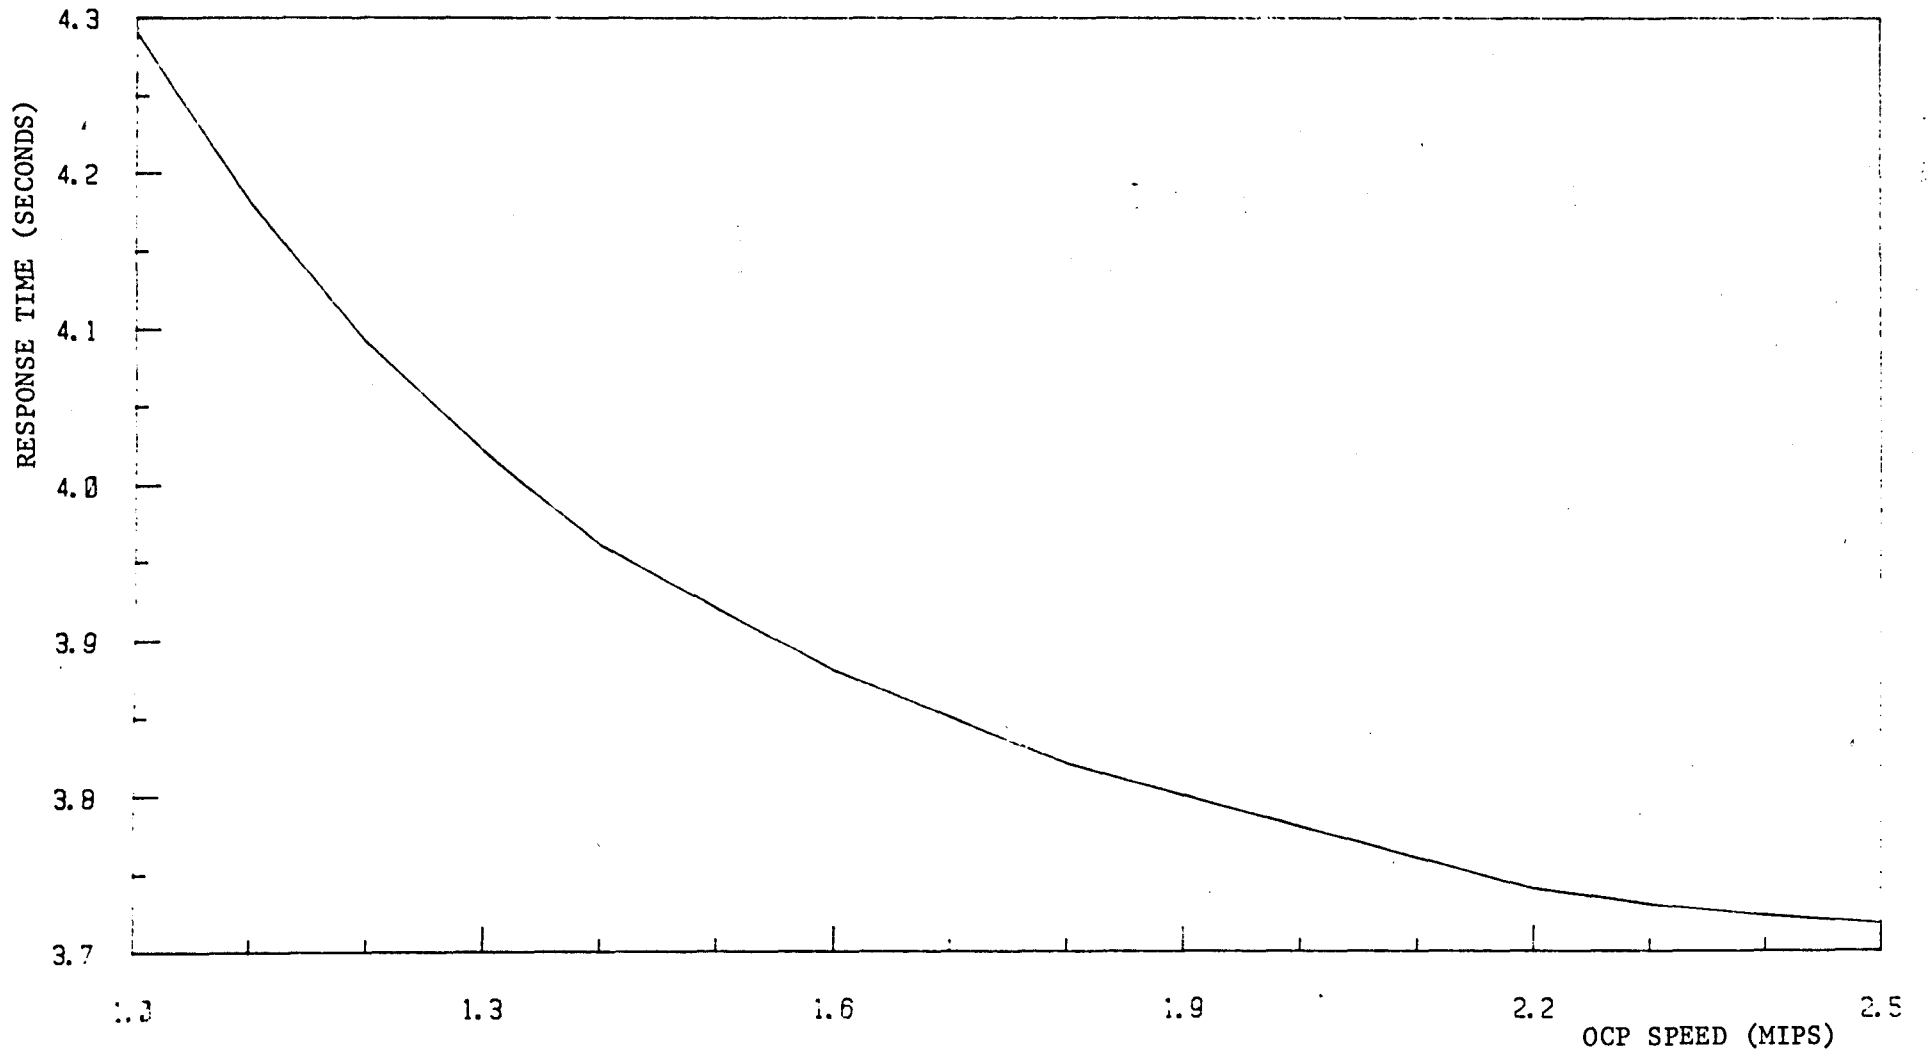
Figure 5.18   Disc traffic equalized. OCP speed vs response time.

6. MEMORY MANAGEMENT MODEL

6.1 INTRODUCTION

Memory management is an important part of a configuration planning exercise. Too little memory will result in excessive page faults being incurred thus causing a possible 'bottleneck' situation at one or more of the disc drives. Too much memory is a financial waste. As explained in section 4.5.1, the main store is not included in the GFSA model, and can therefore not be evaluated using SNAP. Memory management forms a completely separate issue. A special memory model is created to determine the optimum amount of main store required.

The ICL virtual machine concept does not lend itself to an ideal memory management strategy. This chapter first discusses the theory behind memory management and then explains the work done in trying to produce a satisfactory memory model for the GFSA system.

6.2 MEMORY MANAGEMENT THEORY

Under ICL's VME/B operating system, as with all multiprogramming virtual memory systems, many jobs simultaneously compete for OCP processing power and for main memory. While a job is executing, it is allocated a number of page frames greater than or equal to some specified minimum quota size. During its execution, a job may wish to access a page which is not present in main storage. A virtual store interrupt, or VSI, will result, causing the page to be loaded into main memory.

The expected execution time between VSIs or, equivalently, the

average number of instructions per VSI, is obviously a function of the size of main storage and the minimum job store quota. If a job is allocated too many or too few page frames, it tends to use them inefficiently. Idealy, it should be provided with sufficient page frames to maintain its working set in main storage.

Figure 6.1 illustrates the ideal relationship between inter-VSI time and the average main store occupancy. The curve shows a low inter-VSI time for jobs with a small store allocation, a steadily increasing inter-VSI time as the store is increased until the curve starts to taper off (in theory if one gives a job enough store it will eventually load everything it could use and so not require any more VSIs, producing an inter-VSI time of infinity. This is obviously not practical and so the graph is drawn only as far as the upper plateau, although some theorists would prefer it to continue upwards after the plateau).

The idea is to plot the occupancy against the inter-VSI time in an attempt to produce a curve which can be used to determine the optimum amount of main memory. A system with a very disciplined memory allocation strategy should produce a worthwhile curve. This was unfortunately not the case with the GFSA system.

6.3   GFSA MEMORY MODEL INVESTIGATIONS

The problem with the ICL memory policy, so far as the memory model is concerned, is that pages are only discarded when the space is required. A job which is run when the machine is busy will be restricted to its quota size. It is possible that the same job, run in an idle machine, will incur nearly the same number of VSIs, and yet have a reported storage occupancy many times its quota size. This is because pages that are no longer required will remain in store until the space they are occupying is required for some other purpose.
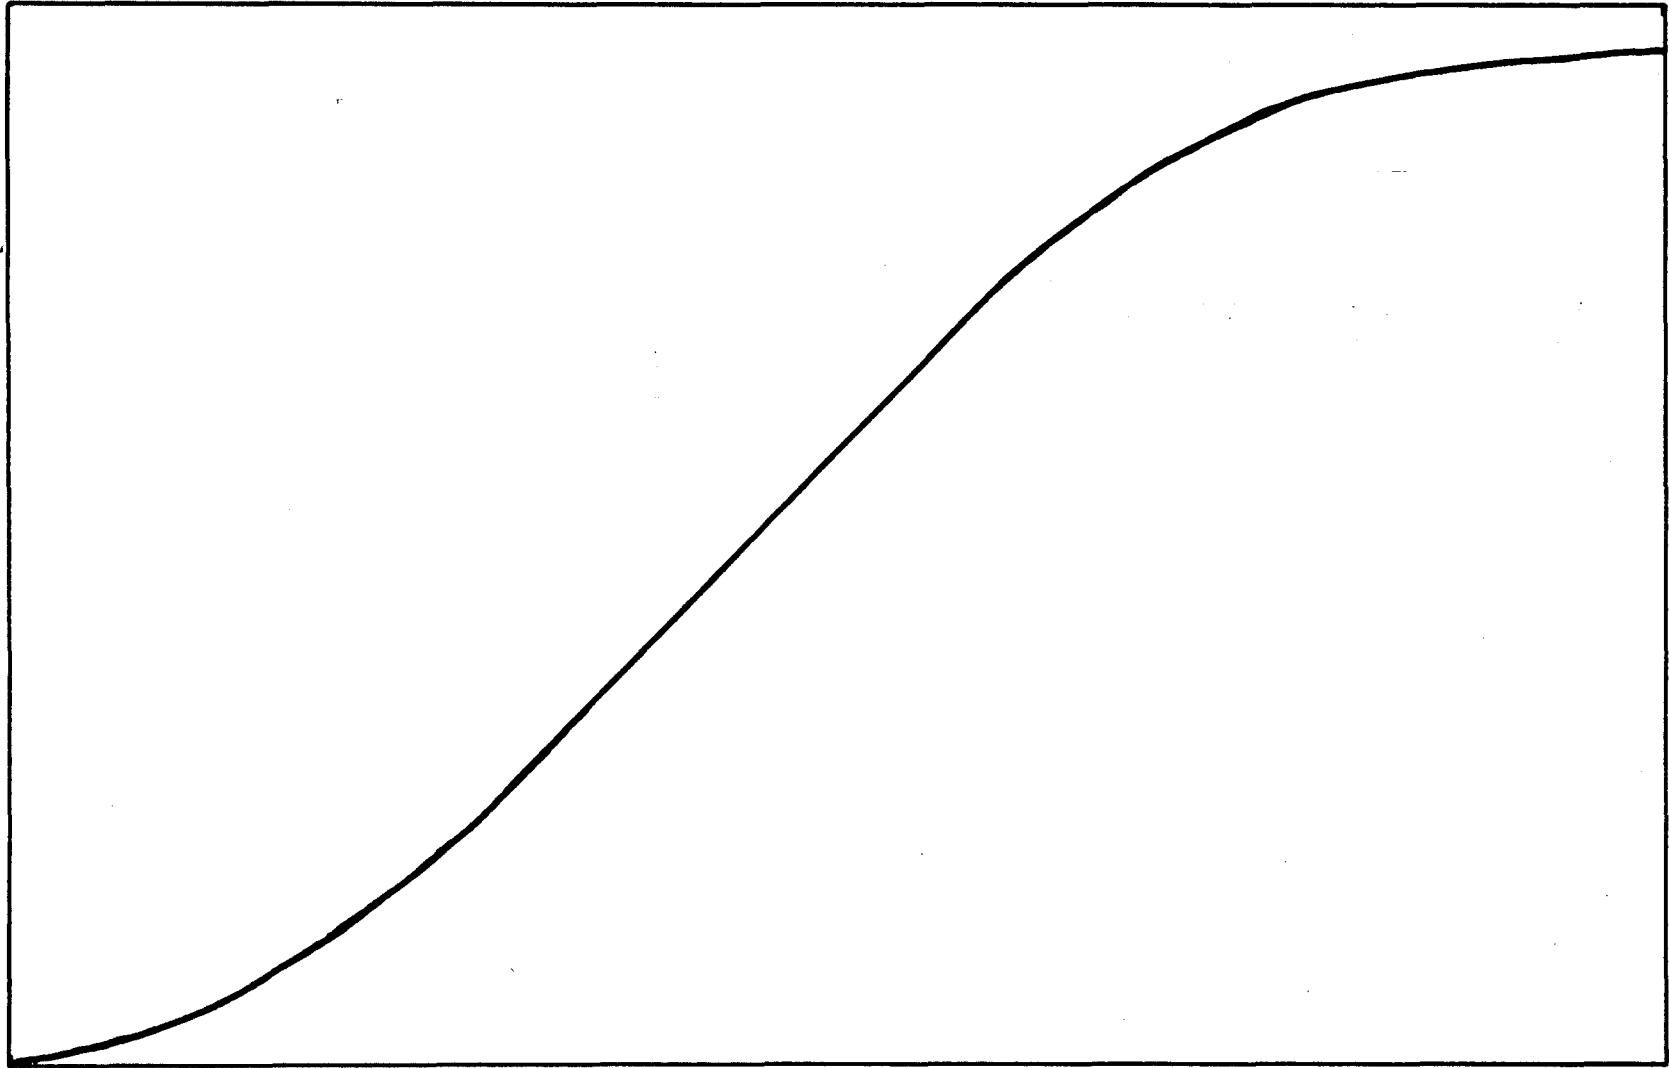
Figure 6.1   Paging behaviour model

The approach taken in attempting to produce a GFSA memory model was as follows. A number of scattergrams were plotted showing main store occupancy versus inter-VSI time. Figures 6.2, 6.4 and 6.6 show the resulting plots for sessions JUN041, JUN081 and JUN182. Figures 6.3, 6.5 and 6.7 are plots of the same sessions, except that the outliers have been removed. Figure 6.8 shows all three sessions together, with the outliers removed.

The first observation is that no curve is apparant. Secondly, the bulk of the points are clustered around the quota size. It is clear that simply producing plots of a number of sessions will not provide a satisfactory curve. A possible strategy would be to continually alter the quota sizes. One could then plot a series of points based on the centre of each cluster. This would mean an enormous amount of processing. The collection of data would have to take place over a long period, as one could not change the quota sizes too frequently. One would also have to be prepared to accept bad performance whilst gathering statistics at either end of the scale - i.e. small quota sizes will incur a VSI overhead, while large quota sizes will incur an unbearable RIRO overhead. Needless to say, this exercise has not been carried out at GFSA.

Figure 6.2   Occupancy - inter-VSI scattergram for JUN041

123

INTER-VSI TIME (msecs)

```
140.0 +
122.5 +        '1    11                                              1

105.0 +        1     111                    1

 87.50 +                          1                              1

 70.00 +        1         1              1
                                                              1
 52.50 +    1         1      1      1 1    1            1           1

 35.00 +                  1         1         1    1
                  1 1   1      1   3 1 1 1    1   1   1 1      1 21
                           1    115    11      1   2      1 11   1   1
 17.50 +        1 1      1       11  1 1 11311  1  2   1 2 1   1 11
                      1   112 326122112211 3   1    1 1 1
                      1  11 1 1 211 1 1       1
                      1        1
  0.000 +
        +....+....+....+....+....+....+....+....+....+....+....+....+....+....+....+....+....+....+
       52.50    87.50    122.5    157.5    192.5    227.5    262.5    297.5    332.5    367.5    402.5    437.5
           70.00    105.0    140.0    175.0    210.0    245.0    280.0    315.0    350.0    385.0    420.0
```

OCCUPANCY (kbytes)

Figure 6.3   Occupancy - inter-VSI scattergram for JUN041 with outliers removed

Figure 6.4   Occupancy - inter-VSI scattergram for JUN081

125
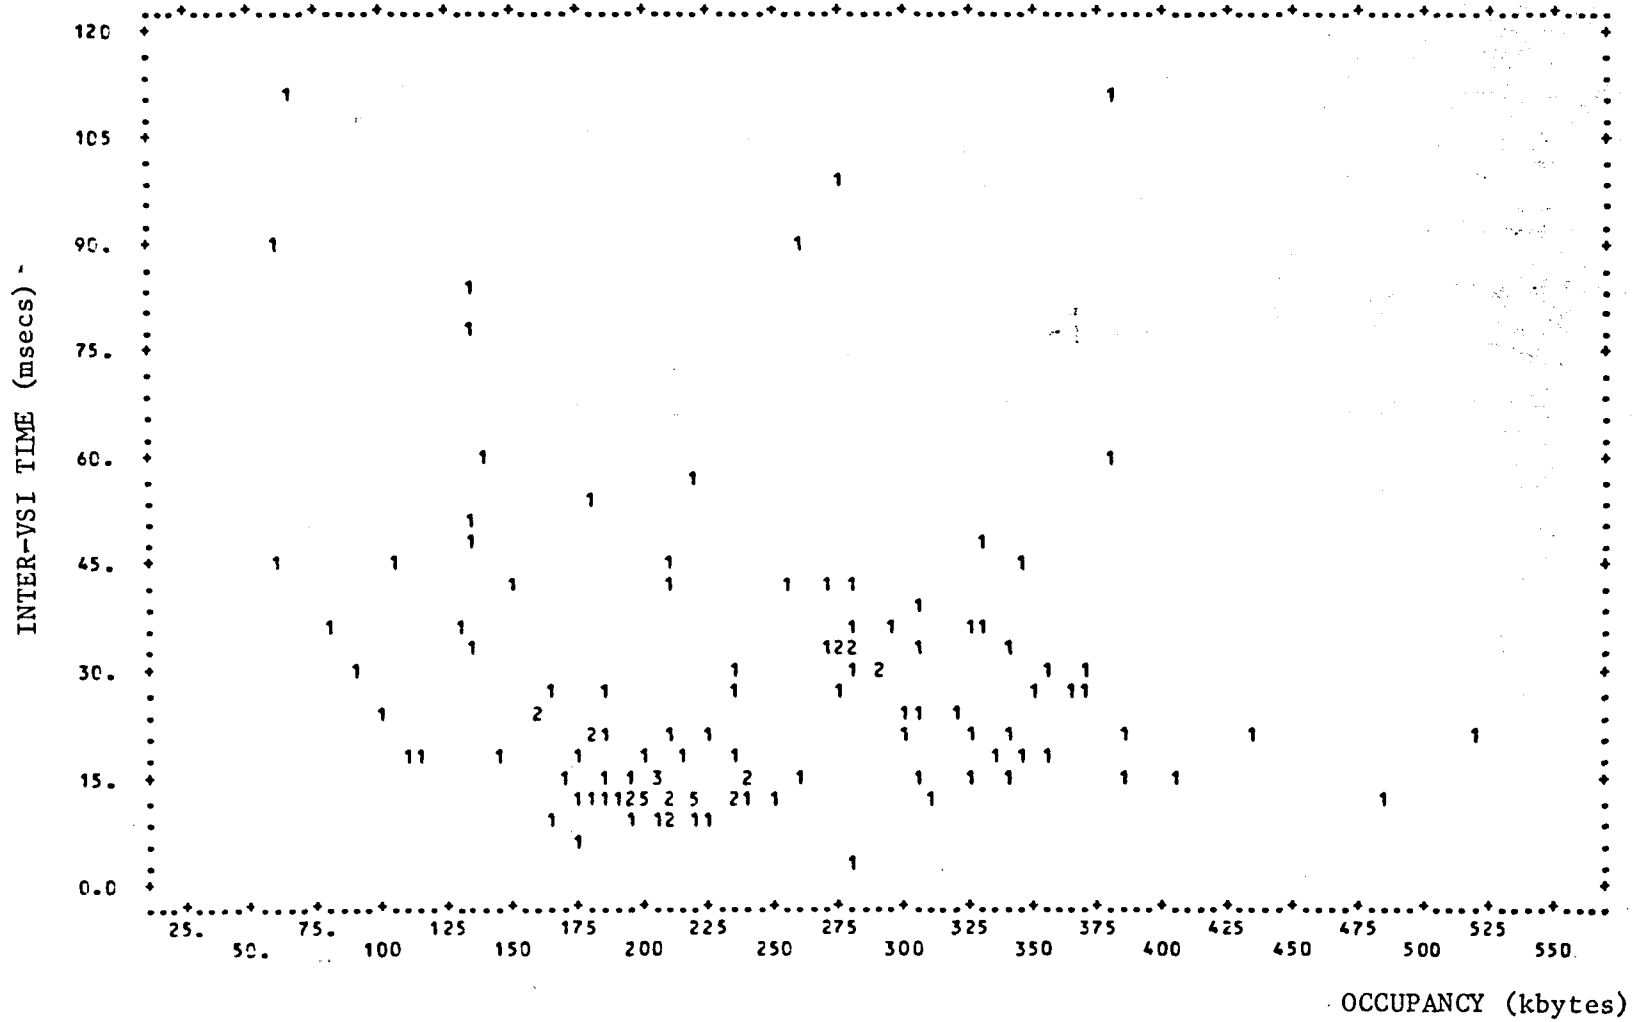
Figure 6.5   Occupancy - inter-VSI scattergram for JUN081 with outliers removed

INTER-VSI TIME (msecs)

```
1000 +  ....+....+....+....+...++...+....+....+....+....+....+....+....+....+....+....+....+....+....+....+....+
       .                            1
       .
       .
 875. +
       .
       .
       .
 75C. +
       .
       .
       .
 625. +
       .
       .
       .
 500. +                          1
       .
       .
       .
 375. +
       .                    1              1
       .
       .
 250. +
       .        11                     1 1        1
       .    1                    1 1   1       1
       .                           1
 125. +    1           1               1                   1
       .      1                1     11         1   1
       .    1  1 1          1    1111 11    1      1
       .      3 2      11  1      1 1 11 2 1 1 3      1
       .     2453B33334483452452523433 12223 111    1      1           1
0.00 +  1 2468677C63211     1      2      1            1
       ....+....+....+....+....+....+....+....+....+....+....+....+....+....+....+....+....+....+....+....+....+
           175     245     315     385     455     525     595     665     735     805     875
         140     210     280     350     420     490     560     630     700     770     840
```

OCCUPANCY (kbytes)

Figure 6.6   Occupancy - inter-VSI scattergram for JUN182
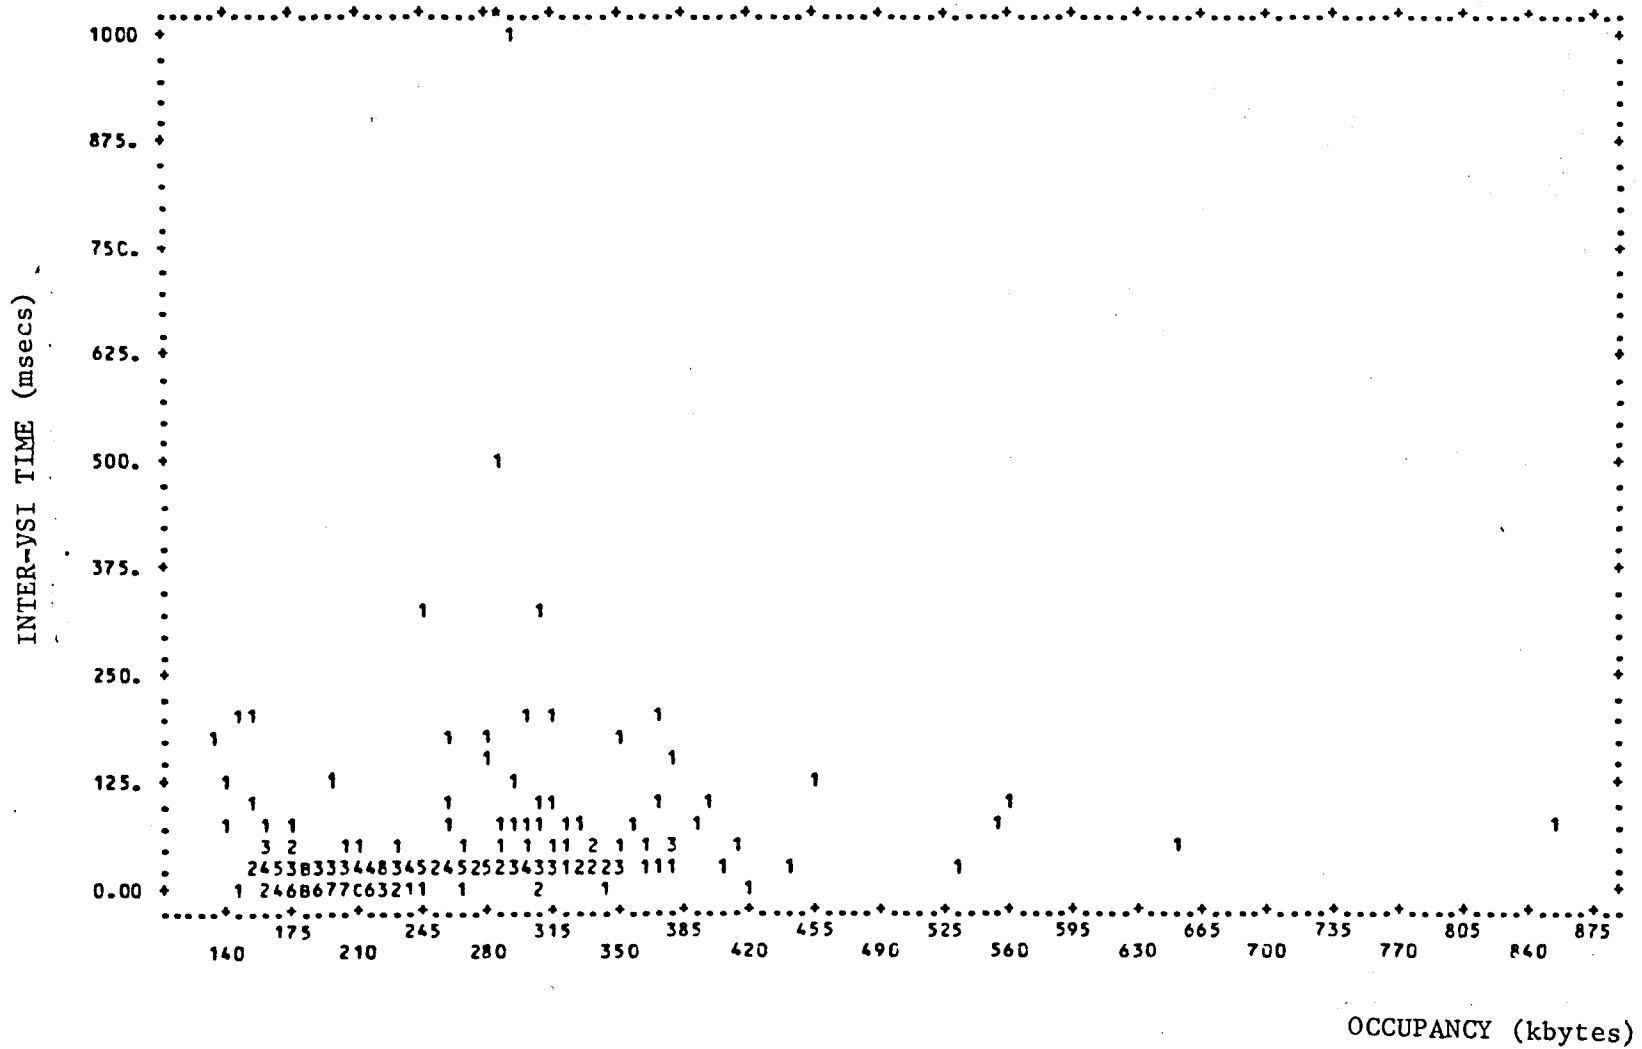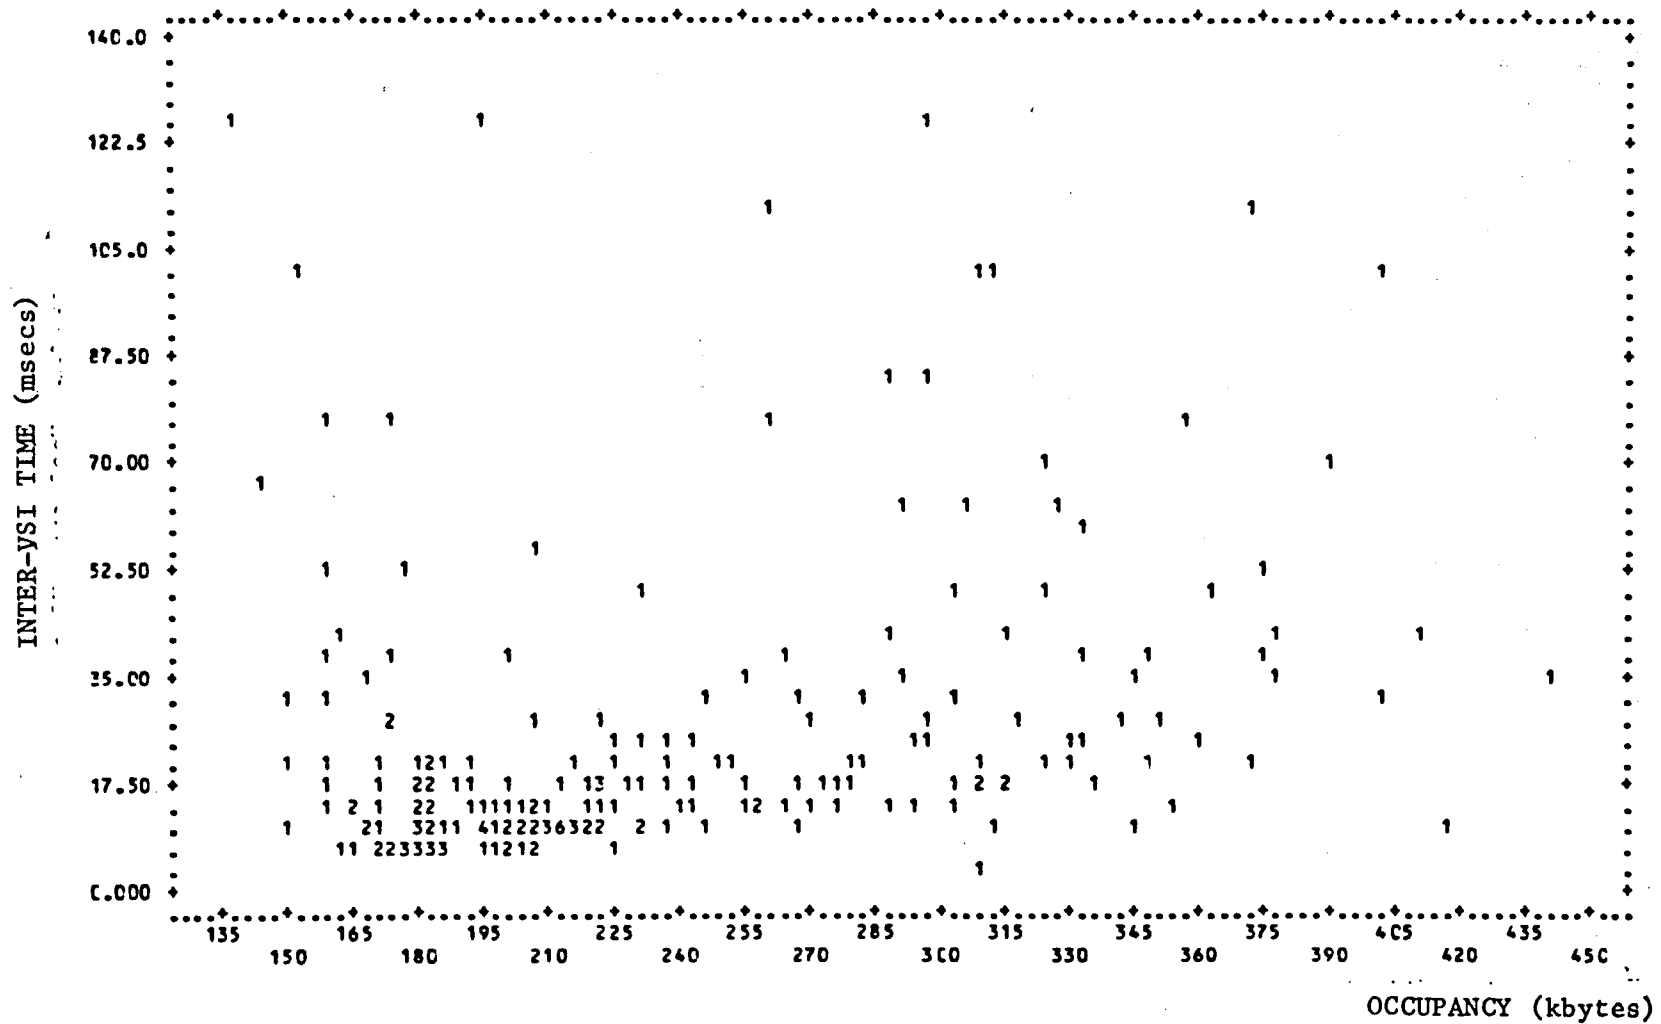
Figure 6.7   Occupancy - inter-VSI scattergram for JUN182 with outliers removed
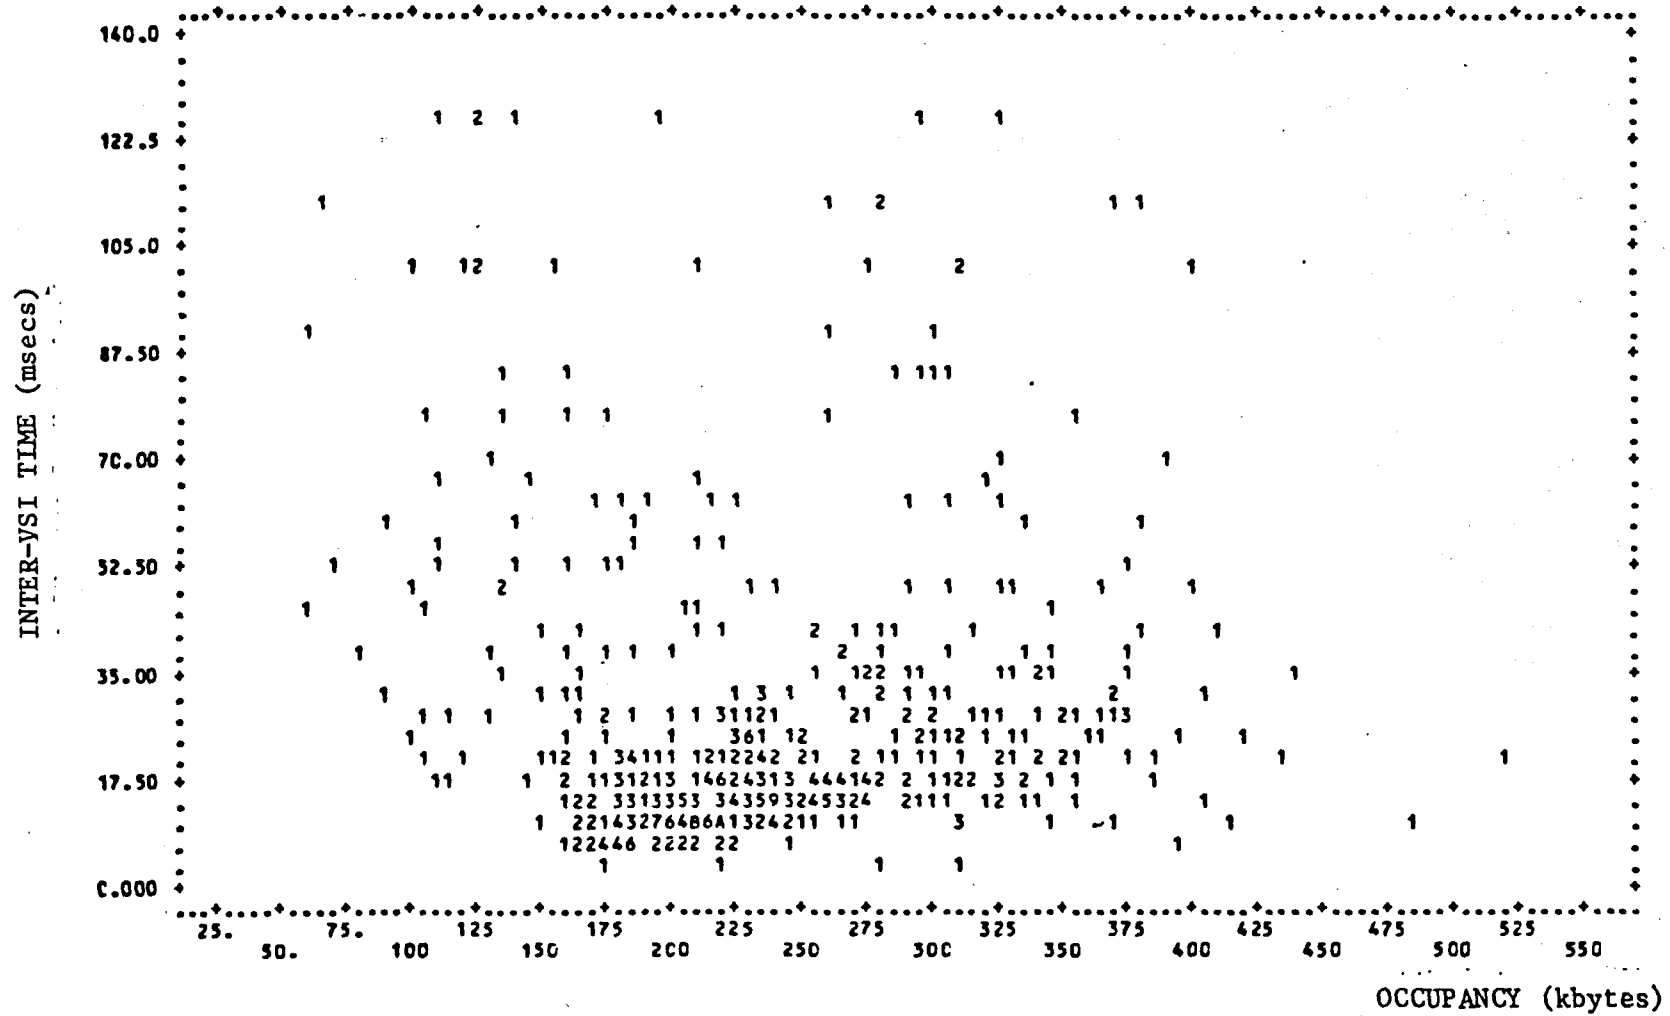
Figure 6.8   Occupancy - inter-VSI scattergram for JUN041, JUN081 and JUN182

APPENDIX A - JOURNAL MESSAGE TYPE (JMT) FORMATS

Only a subset of the possible JMT messages are routed to the performance journal. This appendix gives the format of these JMTs together with an explanation of their use.

The JMT messages are discussed in section 4.2.1. Each message forms one record in the journal, where it is stored in hexadecimal format. Two hexadecimal characters form one byte. A positional character translation is supplied with each JMT format. Note that numerals may either be stored using the character code format (eg. F6F5 = 65), or using a packed hexadecimal format (eg. F6F5 = 63221). The blank spaces in the translation formats are therefore due to numeric hexadecimal values that do not have a character code equivalent.

The message headers follow a standard format. For each message type, byte 8 contains the JMT number. Bytes 9 to 16 contain the time in microseconds since the start of the twentieth century. Bytes 18 to 21 contain the job number of the job that generated the JMT. Not all JMTs contain a PETE number (see section 4.2.1). Where applicable, it is situated in bytes 28 and 29.

Each JMT is dealt with separately on the following pages. In each case the hexadecimal format is given first, followed by the character code translation, followed by an explanation of the purpose of the journal message.

JMT 6 — JOB START MESSAGE

HEXADECIMAL FORMAT

```
00000020D 00000006 000920EE 3D3A818E 03D1F0F0 F9000001 000000EF 9B4DE4E2
D96BE900 075DE2C4 D4D6D5C5 E8404040 40404040 40404040 40404040 40404040
40404040 404DD1D6 C26BE900 0D5DE9F6 F0F3D4D4 C1C3E5F6 F5F5F440 40404040
40404040 40404DD7 D9D66BE9 000A5DD4 C1C3D9C5 E2D7D6D5 C4404040 40404040
40404040 40404040 40404040 40404DC4 C1E36BE2 F1F05DF1 F9F8F161 F0F661F0
F44DE3C9 D46BE2F8 5DF0F97A F0F67AF3 F14DC4C5 E56BE2F4 5DE5C4F9 F4
```

TRANSLATION

```
+                 J009           (US+
+R,Z  )SDMONEY                      +
+     (JOB,Z  )Z603MMACV6554        +
+      (PRO,Z )MACRESPOND           +
+             (DAT,S10)1981/06/0+
+4(TIM,S8)09:06:31(DEV,S4)VD94+
```

EXPLANATION

The JMT 6 Job Start Message is characterized by a PETE number of HEX EF9B.

The following items are extracted from this message:

    Bytes   39 - 69 : Username
    Bytes   79 - 102 : Jobname
    Bytes 112 - 142 : Profile
    Bytes 170 - 177 : Job start time

HEXADECIMAL FORMAT

```
00000405 00000006 000920EE 53A4A06A C3D1F0F0 F5000001 000000EF B44DE4E2
D96BE900 065DE2C4 E3C5E2E3 40404040 40404040 40404040 40404040 40404040
40404040 4C4DD1D6 C26BE900 0D5DE9F7 F1F9E2E3 E2E3C1F3 F2F4F140 40404040
40404040 40404DC4 C1E36BE2 F1F05DF1 F9F8F161 F0F661F0 F44DE3C9 D46BE2F8
5DFCF97A F1F27AF4 F7
```

TRANSLATION

```
+                      J005          (US+
+R,Z  )SDTEST                          +
+      (JOB,Z  )Z719STSTA3241          +
+       (DAT,S10)1981/06/04(TIM,S8+
+)09:12:47+
```

EXPLANATION

The JMT 6 Job End Message is characterized by a PETE number of HEX EFB4.

The job end time (bytes 130 - 137) is the only item extracted from this
JMT. The username and jobname will already have been obtained from the
JMT 6 Job Start Message.

HEXADECIMAL FORMAT

```
00000218 00000009 000920EE 3E3C32FC 02D1FOFO F9201EC3 E4D9D9C5 D5E340E3
D6E3C1D3 E27A202D 40F6F8FO F5F1F8F8 6B404040 40F3F9F2 6B40F2F5 F06140F1
F5F1F2F5 F2F51520 2AC5D3C1 D7E2C5C4 40E5C1D3 E4C5E27A 40F6F8FO F5D26B40
40F3F9F2 6B404040 F040E5E2 40E3D9C1 D5E2C6C5 D9E27A40 40404040 F06140F1
F1F0F340 D4C9D3D3 61E5E2C9 40E3D6E3 C1D37EF6 F4F4F3D2 6140F5F8 F015202A
D27E4040 4040F161 4040F2F2 40D37E40 404040F2 614040F6 F840C47E 40404040
F16140F3 F9F240C7 7E404040 40F16140 4040F840 D97E4040 4040F161 404040F3
40C37E40 404040F1 614040F6 F040E47E 40404040 F7614040 F2F71520 2AC4C9E2
C340E3D9 C1D5E2C6 C5D9E27A 40E47E40 404040F0 40C44DD3 D6C3C1D3 5D7E4040
40F4F040 C44DD7E4 C2D3C9C3 5D7E4040 F3F2F340 D37E4040 40F2F940 E3C1D7C5
40E3D9C1 D5E2C6C5 D9E27E40 404040F0 15202AE2 E8E2E3C5 D440D6E5 C5D9C8C5
C1C4E27A 40F2F6F4 F06140F3 F3F0F6F1 F9F36B40 404040F0 6140F5F8 F2F66B40
40404040 4040F061 40404040 404040F0 6140F4F9 F1F4F7F4 D2614040 40404040
40F0
```

TRANSLATION

```
+                   J009    CURRENT  T+
+OTALS:    6805188,     392,  250/  1+
+512525    ELAPSED VALUES: 6805K,   +
+  392,   0 VS TRANSFERS:      0/  1+
+103 MILL/VSI TOTAL=6443K/  580    +
+K=     1/  22 L=     2/   68 D=    +
+1/  392 G=     1/    8 R=    1/   3+
+  C=     1/  60 U=     7/   27   DIS+
+C TRANSFERS: U=      0 D(LOCAL)=   +
+  40 D(PUBLIC)=  323 L=    29 TAPE+
+  TRANSFERS=     0    SYSTEM OVERHE+
+ADS:  2640/ 3306193,    0/ 5826,  +
+        0/       0/ 491474K/       +
+  0+
```

EXPLANATION

The JMT 9 produces various resource usage figures for the VM that calls it. Some of the values are produced as cumulative totals (i.e. the figure is simply updated each time the JMT 9 is generated). The other values are expressed as elapsed values (i.e. the figure reflects the

amount of the resource used since the previous JMT 9).

Most of the figures given in the summary report of the NRP (see section 4.3) are obtained from the JMT 9. A JMT 9 is logged by calling the macro LOGPERFORMANCEDATA. This macro has a number of parameters which determine which values are to be logged. The JMT 9 format is therefore site dependent and thus the actual field positions will not be detailed here.

JMT 16 – FILE CLOSE DETAILS

HEXADECIMAL FORMAT

```
00000267 00000010 000920EE 41F57C80 03D1F0F0 F9000001 000000F6 384DE2E3
C56BE35D 000101D4 4DE2E3E6 6BE35D00 0101D54D C5D9D96B C95D0000 F6384DD5
E4D46BC9 5D000000 0E4DD7C6 D56BC95D 0BA800B1 4DD9C6C3 6BC95D00 00000004D
D3C6D56B E900075D C9C3D3F9 D3C6F14D D9C6D56B E900095D D3D6C7C9 D5C4C1E3
C14DC2D3 D26BC95D 00000000

00000268 00000010 000920EE 420264CD 03D1F0F0 F9000001 000000F6 394DE2E3
C56BE35D 000101D4 4DE2E3E6 6BE35D00 0101D54D C5D9D96B C95D0000 F6394DD5
E4D46BC9 5D000000 004DD7C6 D56BC95D 0BA800B1 4DD9C6C3 6BC95D00 00000004D
D3C6D56B E900075D C9C3D3F9 D3C6F14D D9C6D56B E900095D D3D6C7C9 D5C4C1E3
C14DC2D3 D26BC95D 00000000

0000026A 00000010 000920EE 4210D430 03D1F0F0 F9000001 000000F6 3A4DE2E3
C56BE35D 000101D4 4DE2E3E6 6BE35D00 0101D54D C5D9D96B C95D0000 F63A4DD5
E4D46BC9 5D000000 014DD7C6 D56BC95D 0BA800B1 4DD9C6C3 6BC95D00 00000004D
D3C6D56B E900075D C9C3D3F9 D3C6F14D D9C6D56B E900095D D3D6C7C9 D5C4C1E3
C14DC2D3 D26BC95D 00000000

0000026F 00000010 000920EE 421CBD28 03D1F0F0 F9000001 000000F6 354DE2E3
C56BE35D 000101D4 4DE2E3E6 6BE35D00 0101D54D C5D9D96B C95D0000 F6354DD5
E4D46BC9 5D000000 004DD7C6 D56BC95D 0BA800B1 4DD9C6C3 6BC95D00 00000004D
D3C6D56B E900075D C9C3D3F9 D3C6F14D D9C6D56B E900095D D3D6C7C9 D5C4C1E3
C14DC2D3 D26BC95D 00000000

00000272 00000010 000920EE 422E43CC 03D1F0F0 F9000001 000000F6 364DE2E3
C56BE35D 000101D4 4DE2E3E6 6BE35D00 0101D54D C5D9D96B C95D0000 F6364DD5
E4D46BC9 5D000000 0E4DD7C6 D56BC95D 0BA800B1 4DD9C6C3 6BC95D00 00000004D
D3C6D56B E900075D C9C3D3F9 D3C6F14D D9C6D56B E900095D D3D6C7C9 D5C4C1E3
C14DC2D3 D26BC95D 00000000

00000275 00000010 000920EE 4233B902 03D1F0F0 F9000001 000000F6 374DE2E3
C56BE35D 000101D4 4DE2E3E6 6BE35D00 0101D54D C5D9D96B C95D0000 F6374DD5
E4D46BC9 5D000000 0B4DD7C6 D56BC95D 0BA800B1 4DD9C6C3 6BC95D00 00000004D
D3C6D56B E900075D C9C3D3F9 D3C6F14D D9C6D56B E900095D D3D6C7C9 D5C4C1E3
C14DC2D3 D26BC95D 00000000
```

TRANSLATION

```
+              5ə  J009          6 (ST+
+E,T)    M(STW,T)    N(ERR,I)   6 (N+
+UM,I)    (PFN,I)     (RFC,I)       (+
+LFN,Z  )ICL9LF1(RFN,Z  )LOGINDAT+
+A(BLK,I)      +
```

```
+                    J009        6 (ST+
+E,T)    M(STW,T)    N(ERR,I)  6 (N+
+UM,I)      (PFN,I)      (RFC,I)      (+
+LFN,Z  )ICL9LF1(RFN,Z  )LOGINDAT+
+A(BLK,I)      +

+                  M  J009        6 (ST+
+E,T)    M(STW,T)    N(ERR,I)  6 (N+
+UM,I)      (PFN,I)      (RFC,I)      (+
+LFN,Z  )ICL9LF1(RFN,Z  )LOGINDAT+
+A(BLK,I)      +

+    ?              J009        6 (ST+
+E,T)    M(STW,T)    N(ERR,I)  6 (N+
+UM,I)      (PFN,I)      (RFC,I)      (+
+LFN,Z  )ICL9LF1(RFN,Z  )LOGINDAT+
+A(BLK,I)      +

+                    J009        6 (ST+
+E,T)    M(STW,T)    N(ERR,I)  6 (N+
+UM,I)      (PFN,I)      (RFC,I)      (+
+LFN,Z  )ICL9LF1(RFN,Z  )LOGINDAT+
+A(BLK,I)      +

+                    J009        6 (ST+
+E,T)    M(STW,T)    N(ERR,I)  6 (N+
+UM,I)      (PFN,I)      (RFC,I)      (+
+LFN,Z  )ICL9LF1(RFN,Z  )LOGINDAT+
+A(BLK,I)      +
```

EXPLANATION

The JMT 16 is generated when a file is closed. More than one JMT 16
is generated for each file, the actual amount being determined by the
type of file. In this example six JMT 16 messages were produced. The
corresponding JMT 62 produced at file open is shown on page A-12.

The SRN is given in bytes 70 to 73 of the message with a PETE number
of HEX F637. This is the last message shown above. The PETE number is
HEX OB = 11. (Note the SRN in the JMT 62 - page A-12).

The actual transfer figure is reported in the same position (bytes 70
to 73) of one or more of the other messages, depending on the type of
file (the reader is referred to the relevant ICL manual for further
details).

JMT 21 - SESSION START MESSAGE

HEXADECIMAL FORMAT

```
00C0000F 00000015 000920EE 21B1BF0E 03D1F0F0 F0000001 000000F0 4F4DE2C9
C46BE900 065DD1E4 D5F0F4F1 40404040 40404040 40404040 40404040 40404DC4
C1E36BE2 F1F05DF1 F9F8F161 F0F661F0 F44DE3C9 D46BE2F8 5DF0F87A F5F87AF4
F1
```

TRANSLATION

```
+                   J000        0!(SI+
+D,Z   )JUN041                      (D+
+AT,S10)1981/06/04(TIM,S8)08:58:4+
+1+
```

EXPLANATION

The JMT 21 start message gives the session identity (bytes 39 to 62),
the session start date (bytes 72 to 81) and the session start time
(bytes 90 to 97).

JMT 21 - SESSION END MESSAGE

HEXADECIMAL FORMAT

```
0000A838 00000015 00092187 6A431F30 03D1F0F0 F0000001 000000F0 504DC4C1
E366E2F1 F05DF1F9 F8F161F0 F661F1F1 4DE3C9D4 6BE2F85D F2F37AF5 F17AF1F7
```

TRANSLATION

```
+                   JOCO         0&(DA+
+T,S10)1981/06/11(TIM,S8)23:51:17+
```

EXPLANATION

The session end date (bytes 39 to 48) and time (bytes 57 to 64) are
provided by this message. Using the session start and session end
messages, the session length can be calculated.

JMT 26 — ACCOUNTING DETAILS

HEXADECIMAL FORMAT

```
00000404 0000001A 000920EE 539FDA7E 03D1F0F0 F5000001 000000F0 184DD9F0
F06BC95D 00000006 4DD9F0F1 6BC95D00 00015B4D D9F1F06B C95D0000 12974DD9
F1F36BC9 5D000000 004DD9F1 F46BC95D 000019CA 4DD9F1F7 6BC95D00 0001174D
D9F2F26B C95D0011 A9E54DD9 F2F56BC9 5D000000 004DD9F5 F56BC95D 00000000
4DD9F1F2 6BC95D00 00014E40 40
```

TRANSLATION

```
+                  = J005      0  (R0+
+0,I)      (R01,I)    $(R10,I)     (R+
+13,I)     (R14,I)     (R17,I)     (+
+R22,I)   V(R25,I)     (R55,I)      +
+(R12,I)    +   +
```

EXPLANATION

The JMT 9 only produces a total OCP time for the job - i.e. this includes the system overhead. The JMT 26 is used to obtain the OCP attributable to the user job only. The JMT 26 is used for accounting details. The format is site dependent and the reader is referred to the relevant ICL manual.

JMT 61 - SYSTEM SNAPSHOT DETAILS

HEXADECIMAL FORMAT

```
000002FB 0000003D 0C0920EE 48449F14 02D1F0F0 F0000000 00000B00 11000000
0C11115C 5CE2E25C 5C00A018 E914215C 5CE2E25C 5C00A019 E414515C 5CE2E25C
5C00B000 721100E2 E8E2F0F0 F100C039 B21110E2 E8E2F0F0 F300F009 951150C7
C6E2C1F1 F300F000 031420E2 E8E2F0F0 F200C001 EF1430C7 C6E2C1F0 F100F000
031440C7 C6E2C1F0 F200F000 5B1450E2 E8E2F0F0 F400F000 061460C7 C6E2C1F0
F300C000 231470C7 C6E2C1F1 F200F000 03FF
```

TRANSLATION

```
+                       J000                +
+      **SS**      7   **SS**     U   **SS*+
+*          SYS001         SYS003  0     &G+
+FSA13  0      SYS002           GFSA01  0   +
+     GFSA02  0  $ &SYS004  0     -GFSA0+
+3         GFSA12  0     +
```

EXPLANATION

The JMT 61 message is logged to the performance journal every two minutes. It records the name of the disc volume mounted on each disc drive.

HEXADECIMAL FORMAT

00000251 0000003E 000920EE 40F2CDF2 02D1FUFC F9000000 0BE2E8E2 F0F0F10
000800

TRANSLATION

+   2 2 J009   SYS001 +
+   +

EXPLANATION

The JMT 62 records the SRN (bytes 22 to 25), the volume name (bytes
26 to 31) and the block size (bytes 32 to 35). The example above has
an SRN of HEX OB = 11 and is in fact the corresponding message for
the JMT 16 on page A-5.

JMT 63 – RIRO INFORMATION

HEXADECIMAL FORMAT

```
000003FD 0000003F 000920EE 53096A3C 03D1F0F0 F5000000 00000000 0F000000
00000000 00000000 00000000 00000000 00000000 0F000000 00000000 00000000
00000000 00000000 0000000C 2A000000 0000000C 27000000 00000000 00000000
000005A8 10000000 00000C8D 2F000000 000078D4 FE000000 00000046 0A000000
00006C96 00000000 00000000 01
```

TRANSLATION

```
+                       J005            +
+                                       +
+                                       +
+                        M              +
+    %             +
```

EXPLANATION

Although the JMT 63 logs many details, it is only used to obtain the
RIRO figure (bytes 22 to 29).

APPENDIX B - NRP REPORTS

The NRP (New Report Program - see section 4.3) produces four reports for each workload class selected. They are:

- Job Summary Report;
- Device Usage Report;
- Statistics Report;
- Concurrency Distribution Report.

Examples of these reports are given here. They are taken from the session of 4 June 1981 (session identity : JUN041) and show the performance data for each of the three workload classes in the GFSA model (i.e. MAC, BATCH and SYSTEM - see section 4.4). Note that the time selection facility has been set to select only those jobs falling in the periods 0858-1255 and 1400-1630.

```
IITTTTTTRRRR
II  TT  RR RR        SESSION IDENTITY JUN041              AT 08:58:41 ON 1981/06/04
II  TT  RRRR
II  TT  RR RR
```

NEW REPORT PROGRAM - SUMMARY REPORT
****

WORKLOAD GROUP USERNAME.JOBNAME.PROFILE	START TIME	ELAPSED TIME (SECS)	OCP SECS USER	OCP SECS TOTAL	M/C SPEED KIPS	VSI /SEC /VSI	USER PAGES XFERS	SYSTEM PAGE XFERS	FILE XFERS	USER DISC XFERS	USER TAPE XFERS	MM XFERS /SEC	DIREC TOR XFERS	OCC (KB)	NO OF RIROS
MASTER.R730SWPRF6571.BATCH	09 08 38	76	3	2	1062	28	2	383	34	1	0	25	51	338	0
CSSTOR.B104STPST6556.BATCH	09 10 10	408	12	11	895	41	1	781	817	973	0	194	1235	418	0
ENGINE.R925SWMEN5071.BATCH	09 13 15	208	7	5	1031	31	1	90	43	1	0	42	215	370	0
MASTER.OPERATOR .BATCH	09 08 51	242	11	10	1035	8	1	638	37	2	0	29	292	109	0
SDINV .R723ISLAC3241.BATCH	09 24 27	444	6	6	995	11	0	613	142	1	0	100	551	302	0
SDMPAY.R714PPCES3041.COMP	09 22 37	287	60	59	978	6	2	1593	399	360	0	18	695	364	0
CSSTOR.B104STN136556.BATCH	09 20 12	1420	46	45	1014	17	1	5950	2071	993	0	200	1609	184	0
MISFB .R601FBJOP6554.BATCH	09 32 01	1156	22	21	977	44	1	6066	2208	2180	0	154	1102	355	0
MASTER.R730SWPRF6171.DEVEL	09 08 52	2035	60	60	807	24	1	7591	1887	1768	20	79	2926	218	2
CSMEDA.B100MAP256554.BATCH	09 44 04	659	12	11	889	36	2	2923	193	175	22	52	374	318	1
SDMPAY.R713MPRJC3141.COMP	09 35 44	656	60	60	946	6	2	3498	419	410	0	23	975	215	0
CCGEOL.R772GBJRS6026.BATCH	09 51 24	1135	23	22	973	49	2	4153	210	172	582	101	1517	270	0
CSPROJ.R723PCLAC3241.BATCH	09 55 11	1673	76	75	893	22	1	8284	354	2600	0	58	1774	204	0
CSMEDA.B100MAP256554.BATCH	10 11 23	732	9	8	920	46	2	4491	217	195	24	69	351	293	1
MISFB .R601FBGRB6554.BATCH	10 33 38	834	20	19	974	46	1	5079	1862	1837	0	148	1043	339	0
TAPES .B8210PMPT5071.BATCH	09 21 29	4827	113	112	893	9	27	26546	9	1	133	200	22293	282	49
SDMPAY.R713MPRJC3141.COMP	10 59 56	338	59	59	942	6	1	838	448	411	0	21	793	236	0
SDMPAY.R713MPRJC3141.COMP	11 07 20	535	43	42	919	15	2	4743	2991	2975	0	89	769	208	0
ACCOUN.B8310PDAL8070.BATCH	10 23 40	2607	44	43	961	37	2	13565	2237	935	1433	92	1629	240	4
MISFB .R601FBGRB6554.BATCH	10 47 37	598	20	19	979	38	0	2522	1755	1730	0	145	1028	373	0
CSMEDA.B100MARE26554.BATCH	10 39 36	2353	24	23	966	60	1	12777	418	288	68	79	1480	301	1
TAPES .B8290PLTD8070.BATCH	11 23 19	179	9	9	878	4	0	1315	5	1	0	70	590	43	0
SYSTEM.R812SWIG15071.COMP	11 21 28	102	4	3	812	20	0	433	307	299	0	126	60	291	0
PERT .Z752PP0156043.MMILL	11 07 01	1859	44	43	973	76	2	16262	711	740	0	93	3288	326	1
CSSTFR.B821TFSHT6550.BATCH	11 40 54	530	17	16	963	9	3	3599	231	1	0	88	1410	282	0
TAPES .Z8190PRUI6770.BATCH	11 55 33	150	4	3	900	12	1	695	5	1	0	18	58	296	0
ACCOUN.B8310PDAL8070.BATCH	11 27 43	1917	49	48	956	39	1	17560	2816	1003	1998	99	1710	309	3
MISFB .R601FBNSE6554.BATCH	11 59 43	451	22	21	971	34	1	941	2378	2353	0	166	1072	377	0
SDTEST.B731STSTC3141.COMP	11 54 22	311	52	51	961	6	1	3332	280	243	0	14	468	177	0
SDMM .R715MMHXX3241.DEVEL	10 46 38	5339	130	129	861	13	1	37135	50917	50461	0	416	3283	107	1
ORORE .B7726DCOR3042.MILL	11 25 30	2260	82	81	904	27	1	19462	282	165	0	51	4008	267	0
SYSTEM.R812SWIG35071.COMP	12 19 32	17	2	1	1028	36	0	36	32	1	0	43	22	661	0
TAPES .Z8190PRUI6770.BATCH	12 13 04	38	4	3	893	9	0	251	32	1	0	19	60	506	0
MASTER.R812SWIG25071.BATCH	11 59 40	1635	124	122	523	3	2	7641	4977	1961	3005	81	4897	25	0
SDMM .R727MMB393241.BATCH	12 27 11	49	1	0	1054	66	0	35	5	1	0	48	16	319	0
SDTEST.B731STMCN3041.BATCH	12 04 33	1036	55	54	579	8	3	5239	260	199	641	60	2388	327	2
SDMM .R727MMB393241.BATCH	12 28 12	50	1	0	1029	57	0	82	5	1	0	49	16	324	0
SDMM .R727MMR393241.BATCH	12 29 19	39	1	0	1034	60	0	278	5	1	0	46	15	326	0
SDMM .R727MMB393241.BATCH	12 30 09	77	1	0	1028	60	0	222	5	1	0	51	17	159	0
MISFB .R601FBJOP6554.BATCH	12 07 18	608	22	21	981	37	0	3425	2206	2178	0	156	----	---	-
MASTER.R730SWGET5071.DEVEL	09 08 51	12731	380	378	1040	2	46	78738	40	3	0	1			
SYSTEM.R812SWIG45071.COMP	12 22 35	322	26	25	984	4	2	1444	9	1	0	17			
ACCOUN.B8310PCOP8070.BATCH	12 31 37	1120	25	24	864	31	2	9252	4171	688	3463	-200			
MISFB .R601FBJOP6554.BATCH	14 11 51	431	23	22	991	36	1	1000	2359	2304	0	158			
CSMEDA.R822SWCSN5C71.BATCH	14 17 29	327	33	32	839	16	2	2000	3227	3171	3182	217			
SYSTEM.R817SWKEN5071.BATCH	14 57 03	28	2	1	1075	23	0	0	32	1	0	82			
CCGEOL.R822SWCSN5C71.BATCH	14 20 35	1863	148	147	776	4	6	8000	17858	17826	17873	248			
SDTEST.B731STMCN3041.BATCH	14 33 09	1268	119	117	441	3	3	6000	276	223	293	54			
MISFB .R601FBNSE6554.BATCH	14 30 07	799	23	21	952	34	1	5000	4333	4272	0	184			
```
B-2
```

BATCH WORKLOAD GROUP

SUMMARY REPORT

PAGE 1

```
 SDTEST.R731STSRR3141.COMP      15 04 12    206    14    13 1046   17    2      0     322    280      0   145   1625   459   0
 SDPLAN.R822SWCSN5C71.BATCH     14 56 26    649    33    32  739   15    3   3C00    4181   4126   4137   270    417   110   1
 CSPROJ.R822SWCSN5071.BATCH     14 58 21    533    22    22  820   20    2   3000    2094   2066   2077   212    427   238   1
 TAPES .B8210PNPT5070.BATCH     15 10 57    514    26    25  894    5   29   2000       8      1     28   159   4021   192   7
 SDTEST.R731STSRR3141.COMP      15 10 03    492    62    61  976    6    1   2000     824    801      0    69   3446   219   0
 SDTEST.R719STSTA3241.DEVEL     15 11 47    204     5     4 1034   83    1   1000     105     60      0    90    328   523   0
 SDMPAV.R713MPRJC3141.BATCH     15 17 23    251    42    41  922    8    0      0    2991   2975      0    91    745   124   0
 ORORE .R7666TRHW3242.BMILL     14 36 09   3985   912   911 1002    3    1  18C00    1451   2975      0     9   5099    26   0
 CSACCT.B100AC3116554.BATCH     15 51 08    637    18    17  869   40    1   5000    1339   1533      0   162   1302   395   0
 CSINV .R723ISLAC3241.BATCH     16 01 59    392    27    26  908   21    1   4000     245    198      0    27    515   401   0
 MASTER.R812SWIGK5071.COMP      15 25 08   2110   144   142  589    2    2  13000   11850  11807      0   133   7008    22   0
 SDMM .R715PMHTM3261.DEVEL      15 56 12    144     5     4  977   44    1   1000     322    287      0   125    230   351   0
 SDINV .R723ISLAC3241.BATCH     16 08 36    339    10     9 1003   12    0   1000      72      1      0   194   1747   303   0
**********************************************************************************************************************************
        WORKLOAD GROUP          * START  *ELAPSED* OCP SECS  * M/C * VSI *USER *SYSTEM* FILE * USER * USER *  MM  *DIREC * OCC *NO OF
 USERNAME.JOBNAME.PROFILE        *  TIME  * TIME  ***************SPEED* /SEC*PAGES* PAGE * XFERS* DISC * TAPE * XFERS*TOR   *(KB) *RIROS
                                 *        *(SECS) *USER *TOTAL*KIPS *     */VSI *XFERS *      * XFERS* XFERS* /SEC *XFERS *     *
**********************************************************************************************************************************
```

THERE ARE   62 JOBS IN THIS WORKLOAD GROUP

WORKLOAD DEFINITION
****

PROFILE=BATCH&COMP&DEVEL&MILL&BMILL
TIME=085800&125500&081400000&163000&10

SESSION END AT 01:05:54 ON 1981/06/05 ELAPSED TIME 58024 SECONDS WITH A TOTAL OF 299 JOBS IN THIS SESSION

BATCH WORKLOAD GROUP

SUMMARY REPORT

```
IITTTTTTRRRR
II  TT  RR RR              SESSION IDENTITY JUN041                    AT 08:58:41 ON 1981/06/04
II  TT  RRRR
II  TT  RR RR
```

NEW REPORT PROGRAM - DEVICE USAGE REPORT
****

```
**************************************************************************************************
* DEVICE   * VOLUME *    BLOCKSIZE    *  FILE   * FILE TRANSFERS * LOADER ACCESSES * OTHER ACCESSES *
* ADDRESS  *   ID   ********************* ACCESSES *******************************************************
*          *        *  AVE  *  VAR   *         * TOTAL * AS %  * TOTAL * AS % * TOTAL * AS % *
**************************************************************************************************
* 01010100  SYS003    2048        0      182    22823  14.382      0    0.000      0   0.000*
* 01010000  SYS001     157   297989       73     5135   3.670   1682   87.015      0   0.000*
* 01040200  SYS002    2048        0        1     9691   6.925     62    3.207      0   0.000*
* 01040700  GFSA12    2048        0       29    41201   1.249      0    0.000      2   0.444*
* 01040400  GFSA02     524   800214      136    16597  11.861     99    5.122    405  90.000*
* 01010500  GFSA13    3792 16128949       20     5283   3.775      0    0.000     29   6.444*
* 01040600  GFSA03    1804   440566       30      679   0.485     12    0.621      0   0.000*
* 01050000  DEV214    4096  4194304        2      279   0.199      0    0.000      0   0.000*
* 01050001  PRF019    6144        0        1        3   0.002      0    0.000      0   0.000*
* 01050002  DEV160    2048        0        4      567   0.405      0    0.000      0   0.000*
* 01040300  GFSA01    1093  1045786       83     1541   1.101     78    4.035      0   0.000*
* ********            2783  2527630        7     1665   1.190      0    0.000     10   2.222*
* 01050103  DEV055    2048        0        4     2074   1.482      0    0.000      0   0.000*
* 01050105  OFF006    3802  3117361        3    21008  15.013      0    0.000      0   0.000*
* 01050104  OFF008    2329   236883        4     3302   2.360      0    0.000      0   0.000*
* 01050101  DEV133    3413  3730545        3     3964   2.833      0    0.000      0   0.000*
* 01010600  GF0053    6144        0        0        0   0.000      0    0.000      4   0.889*
* 01050100  OFF007    3172        0        1     4122   2.946      0    0.000      0   0.000*
**************************************************************************************************
* DEVICE   * VOLUME *    BLOCKSIZE    *  FILE   * FILE TRANSFERS * LOADER ACCESSES * OTHER ACCESSES *
* ADDRESS  *   ID   ********************* ACCESSES *******************************************************
*          *        *  AVE  *  VAR   *         * TOTAL * AS %  * TOTAL * AS % * TOTAL * AS % *
**************************************************************************************************
```

SESSION END AT 01:05:54 ON 1981/06/05 ELAPSED TIME 58024 SECONDS WITH A TOTAL OF 299 JOBS IN THIS SESSION

BATCH WORKLOAD GROUP

DEVICE USAGE REPORT

```
IITTTTTTRRRR
II  TT  RR RR        SESSION IDENTITY JUNQ41              AT 08:58:41 ON 1981/06/04
II  TT  RRRR
II  TT  RR RR
```

STATISTICS FOR THIS WORKLOAD GROUP
****

		TOTAL	AVERAGE	VARIANCE
OCP TIME(SECS)	TOTAL	3404.846	54.92	15220.68
	USER	3461.232	55.83	15247.66
OCP UTILISATION			16.51	0.23
VIRTUAL STORE INTERRUPTS	TOTAL	38330	618.23	432541.92
	USER	40834	658.61	433466.13
INTER VSI TIME(MSECS)			88.83	13908.63
OCCUPANCY(KB)			280.63	18013.66
VIRTUAL STORE PAGES TRANSFERED	DISC TOTAL	109362	1763.90	23972897.03
	USER	110598	1783.84	24031094.15
	DRUM TOTAL	0	0.00	0.00
	USER	0	0.00	0.00
DIRECTOR TRANSFERS	LOCAL	5372	86.65	25774.27
	PUBLIC	59898	966.10	7968495.48
	LOWER	40400	651.61	1408360.50
	TOTAL	105670	1704.35	9240962.76
I/O-TRANSFERS	DISC	133218	2148.68	45851174.94
	TAPE	38979	628.69	5716451.63
	FILES	142317	2295.44	6117170320.24
MM-TRANSFERS	TOTAL	284213	4584.08	73319673.31
	USER	286685	4623.95	73309197.02
NO OF FILES ACCESSED		786	12.68	97.84
NO OF RECORDS SPOOLED		0	0.00	0.00
NO OF RIROS		132	2.13	73.37
ELAPSED TIME		69215	1116.37	3398461.08

SESSION END AT 01:05:54 ON 1981/06/05 ELAPSED TIME 58024 SECONDS WITH A TOTAL OF 299 JOBS IN THIS SESSION

BATCH WORKLOAD GROUP

STATISTICS REPORT

```
IITTTTTTRRRR
II  TT  RR RR              SESSION IDENTITY JUN041                    AT 08:58:41 ON 1981/06/04
II  TT  RRRR
II  TT  RR RR
```

```
                                                    TOTAL TIME(SECS)   NORMALISED
                        CONCURRENCY DISTRIBUTION
                        -------------------------
                                  LEVEL  0          4912            23.81
                                  LEVEL  1          3477            16.86
                                  LEVEL  2          3258            15.79
                                  LEVEL  3          2570            12.46
                                  LEVEL  4          5850            28.36
                                  LEVEL  5          4680            22.69
                                  LEVEL  6           791             3.84
                        AVERAGE CONCURRENCY          3.35
                        -------------------------
```

SESSION END AT 01:05:54 ON 1981/06/05 ELAPSED TIME 58024 SECONDS WITH A TOTAL OF 299 JOBS IN THIS SESSION

BATCH WORKLOAD GROUP

CONCURRENCY DISTRIBUTION REPORT

INSTITUUT VIR TOEGEPASTE REKENAARWETENSKAP/INSTITUTE FOR APPLIED COMPUTER SCIENCE - UNIVERSITY OF STELLENBOSCH - SOUTH AFRICA

```
IITTTTTTRRRR
II  TT  RR RR       SESSION IDENTITY JUN041              AT 08:58:41 ON 1981/06/04
II  TT  RRRR
II  TT  RR RR
```

NEW REPORT PROGRAM - SUMMARY REPORT
****

WORKLOAD GROUP USERNAME.JOBNAME.PROFILE	START TIME	ELAPSED TIME (SECS) USER	OCP SECS USER TOTAL	M/C SPEED KIPS	VSI /SEC	USER PAGES /VSI	SYSTEM PAGE XFERS	FILE XFERS USER	USER DISC XFERS	USER TAPE XFERS	MM XFERS /SEC	DIREC TOR XFERS	OCC (KB)	NO OF RIROS
S DTEST.Z719STSTA3241.MACRESPOND	09 08 37	250	5	4 1053	66	24	840	58	19	0	62	234	237	15
SYSTEM.Z818SWKJF5071.MACRESPOND	09 12 05	89	3	2 1094	55	11	60	33	22	0	68	89	289	2
S DMONE.Z603MMACV6554.MACRESPOND	09 06 31	471	7	7 1047	80	9	1588	68	25	0	48	297	224	11
S DMPAY.Z714MPCES3041.MACRESPOND	09 11 01	522	13	12 984	48	11	1424	441	377	0	101	856	255	15
S DMONE.Z723MMLAC6054.MACRESPOND	09 18 11	184	5	5 994	57	17	373	58	20	0	73	341	274	10
S MC .Z818SWKJF5071.MACRESPOND	09 11 23	941	12	11 1120	57	35	3737	114	151	0	70	609	218	49
S DINV .Z723ISLAC3241.MACRESPOND	09 21 39	140	3	2 993	84	3	648	22	14	0	62	125	254	1
INVGOL.Z630INMJG6052.MACRESPOND	09 15 17	927	12	12 1000	55	19	3708	256	132	0	48	430	260	27
C SMEDA.Z729MADKP3241.MACRESPOND	09 18 34	1019	15	14 1000	47	80	3613	103	79	0	55	715	220	130
S DMPAY.Z713MPRJC3041.MACRESPOND	09 10 55	1535	10	9 997	52	81	5123	265	258	0	78	458	230	89
SYSTEM.Z818SWKJF5071.MACRESPOND	09 41 42	285	7	6 1040	40	11	1530	15	0	0	39	240	230	6
CSCAP .Z724CXSRP3241.MACRESPOND	09 31 40	851	14	13 1015	41	10	4161	862	814	0	117	753	174	12
SYSTEM.Z817SWKEN5071.MACRESPOND	09 31 39	915	9	8 1036	39	41	4611	178	164	0	85	529	225	30
ORROCK.Z771RMMAT3042.MACRESPOND	09 43 69	246	5	4 990	118	30	1164	26	58	0	57	170	244	33
CSPROJ.Z723PCLAC3241.MACRESPOND	09 25 13	1743	68	67 961	16	27	7444	828	986	0	35	1405	188	66
C CGEOL.Z7726BJBS6026.MACRESPOND	09 40 49	854	12	12 1129	20	21	5050	23	14	0	17	190	102	11
ORCONV.Z764066DK3042.MACRESPOND	09 45 52	888	16	15 1014	87	6	3998	131	69	0	74	1062	229	29
SYSTEM.Z822SWCSN5071.MACRESPOND	09 14 50	3054	20	19 1046	113	8	12922	105	61	0	60	1107	208	54
C SMEDA.Z729MADKP3241.MACRESPOND	10 01 20	327	8	8 1059	31	26	683	25	120	0	46	240	234	14
O RORE .Z7666TRHW3242.MACRESPOND	09 50 06	1003	143	142 993	9	24	3232	246	401	0	11	1199	43	66
O RANB .Z7680GREC3042.MACRESPOND	09 55 26	708	4	3 979	72	85	2081	28	15	0	81	248	266	46
SYSTEM.Z818SWKJF5071.MACRESPOND	09 46 50	1355	38	37 1032	15	61	5484	36	6	0	14	520	82	76
S DMM .Z727MMB393241.MACRESPOND	09 16 36	3567	96	96 938	34	12	16052	8931	9202	0	136	3846	302	87
ENGINE.Z925HWHHH6769.MACRESPOND	09 54 51	1850	53	51 1104	40	4	9312	1418	1966	0	90	2622	251	16
C SMORE.Z728MOIAN3141.MACRESPOND	09 16 39	4325	21	20 990	42	41	21260	189	121	0	42	720	284	81
SYSTEM.Z817SWKEN5071.MACRESPOND	10 30 04	145	3	2 1074	75	9	1068	8	4	0	70	133	205	3
MASTER.Z812SWIGK5071.MACRESPOND	10 11 19	1542	111	109 1083	6	99	10943	17	1	0	3	334	70	153
S DMM .Z718MMWHI3041.MACRESPOND	09 43 28	3291	13	12 1065	54	32	18699	125	79	0	62	692	217	51
S DMM .Z715MMHXX3241.MACRESPOND	10 07 15	1947	46	45 989	31	28	13842	213	178	0	23	837	235	87
C SMORE.Z724MOMRA3041.MACRESPOND	09 46 47	3182	147	146 1137	12	18	18834	1105	1054	0	16	1355	87	70
SUPPOR.Z817SWKEN5071.MACRESPOND	09 51 02	3221	82	81 1049	7	24	19852	34	1606	1587	249	17082	70	33
SUPPOR.Z811SWIAN5071.MACRESPOND	10 37 02	495	59	59 1015	7	76	4073	8	0	0	2	95	36	69
PERT .Z752PP0156043.MACRESPOND	09 10 05	5778	62	61 1007	70	17	30491	775	767	0	120	6549	277	161
TAPES .Z8190PRUI6770.MACRESPOND	10 48 55	408	3	2 995	80	14	2177	21	13	0	56	119	276	6
O RSTAT.Z765SWGDK3342.MACRESPOND	09 59 44	3969	12	12 1038	113	6	22070	98	53	0	57	819	249	31
SYSTEM.Z812SWIGK5071.MACRESPOND	10 39 38	1599	63	62 1003	8	136	6713	1465	1434	0	34	645	125	154
SUPPOR.Z811SWIAN5071.MACRESPOND	10 56 06	543	3	2 1093	54	30	1195	9	0	0	76	169	219	8
O RORE .Z7666TRHW3042.MACRESPOND	11 06 13	304	6	6 974	49	31	163R	25	14	0	62	348	235	21
S DTEST.Z731STMCN3041.MACRESPOND	10 39 26	1985	24	23 1013	66	9	9367	129	93	0	53	1143	231	30
O RORE .Z7726BCOR3042.MACRESPOND	10 16 07	3432	29	29 1039	43	35	20380	563	684	0	71			
S DMPAY.Z713MPRJC3041.MACRESPOND	10 17 29	3875	17	16 969	51	109	23966	537	484	0	85			
S DMPAY.Z714MPCES3041.MACRESPOND	10 44 25	2268	31	31 932	43	8	11845	129	46	30	49			
S DMM .Z727MMB393241.MACRESPOND	11 10 23	872	23	22 959	57	5	7769	1592	1529	0	129			
SUPPOR.Z811SWIAN5071.MACRESPOND	11 23 15	231	23	23 1013	5	102	1963	6	0	0	2			
INVGOL.Z723INRMH6052.MACRESPOND	11 27 57	2R9	12	12 1009	37	8	1777	222	130	0	50			
TAPES .Z821SWTAP5C71.MACRESPOND	11 15 12	122	4	4 992	53	15	925	50	17	0	47			
CSCAP .Z723CXLAC6041.MACRESPOND	11 07 01	1645	36	36 994	55	14	14445	423	20	0	105			
SUPPOR.Z811SWIAN5071.MACRESPOND	11 33 57	551	13	12 911	17	38	3743	20	16	0	65			
S DMPAY.Z713MPRJC3041. MACRESPOND	11 38 26	998	14	14 972	40	87	7368	494	446	0	37			

MACRESPOND WORKLOAD GROUP

SUMMARY REPORT

PAGE 1

ENGINE.ZY25HWENG6769.MACRESPOND	11	12	03	1311	67	65	793	10	7	10776	47	16	12650	213	1089	99	10
SUPPOR.Z822SWCSN5071.MACRESPOND	10	10	32	5851	54	53	1093	83	7	39871	155	75	0	67	3473	163	122
CSINVM.Z100INMJ66052.MACRESPOND	10	42	42	4278	46	45	1013	46	8	29750	287	224	0	37	1436	231	32
CCROCK.Z827RMFCN3042.MACRESPOND	10	28	30	5270	16	15	922	127	271	39140	90	221	0	61	710	218	1242
SDBPTE.Z728APDKP3241.MACRESPOND	11	13	23	2649	33	32	1012	92	7	23754	727	749	0	69	1503	220	54
MASTER.Z812SWIGK5071.MACRESPOND	11	43	33	897	28	28	1134	19	54	9642	112	99	0	23	539	72	63
SDMM .Z718MMI903041.MACRESPOND	11	35	12	1008	8	8	1048	53	34	8260	56	31	0	59	420	211	32
MPLAN .Z790MLWES6044.MACRESPOND	10	13	03	6330	169	167	845	36	10	46649	1884	1671	0	36	4330	201	131
SDTEST.Z731STMCN3041.MACRESPOND	11	45	48	821	12	11	1037	59	11	9147	100	81	0	61	578	289	16
SDBPTE.Z72RBPDKP3241.MACRESPOND	11	59	04	231	7	7	951	50	12	772	75	66	0	54	298	270	9
SDISD .Z723SWLAC3241.MACRESPOND	11	35	00	2066	54	54	980	14	14	15429	351	3001	0	150	5034	128	24
ORANB .Z7680GREC3042.MACRESPOND	11	14	10	3115	50	49	921	25	21	25585	2951	3587	0	108	1742	187	60
ORORE .Z7766GTRHW3042.MACRESPOND	10	08	04	7410	232	231	987	2	13	52244	139	121	0	45	10279	16	15
SYSTEM.Z818SWKJF5071.MACRESPOND	12	03	15	600	4	3	1038	55	77	2983	8	2	0	50	141	216	28
PERT .Z752PP0156043.MACRESPOND	11	51	57	1176	5	4	983	74	59	7962	53	16	0	101	437	222	45
SDMM .Z727MMB393241.MACRESPOND	11	54	37	1177	24	23	979	38	15	7327	213	231	0	41	740	272	29
SUPPOR.Z822SWPRI5071.MACRESPOND	12	11	19	312	5	4	997	72	15	1431	28	17	0	88	316	232	10
SYSTEM.Z812SWIGK5071.MACRESPOND	12	11	15	582	25	24	1076	12	75	2557	21	4	0	12	286	162	50
SYSTEM.Z822SWCSN5071.MACRESPOND	12	17	46	245	5	4	1070	69	18	1435	44	26	0	82	294	237	11
CSCAP .Z108CXFEW6080.MACRESPOND	09	16	35	11233	34	33	961	52	85	69782	2355	2289	0	112	1399	188	366
ORORE .Z766GTRHW3042.MACRESPOND	12	13	39	678	14	13	1074	109	7	3595	90	59	0	60	704	240	42
TAPES .Z8190PRUI6770.MACRESPOND	11	49	03	2226	27	27	884	16	22	15707	28	15	14	217	5767	169	22
SDTEST.Z711INDDD3041.MACRESPOND	12	22	38	267	4	4	1006	87	24	1666	25	57	0	69	189	234	17
INVGOL.Z711INDDD3041.MACRESPOND	12	28	00	101	3	3	1001	111	6	975	22	13	0	57	147	218	5
ORANB .Z7680GREC3042.MACRESPOND	12	27	50	320	5	4	990	70	9	2831	68	97	0	78	211	231	5
SDMPAY.Z713MPRJC3041.MACRESPOND	12	21	39	560	8	8	997	48	10	4348	99	118	0	102	678	229	8
SUPPOR.Z817SWKEN5071.MACRESPOND	11	58	07	2113	7	6	1024	60	25	12157	17	1	0	102	659	219	23
ORCONV.Z7640GGDK3042.MACRESPOND	10	45	56	6509	19	19	969	256	11	44671	170	111	0	59	1002	220	314
SYSTEM.Z818SWKJF5071.MACRESPOND	12	26	30	538	8	7	1035	89	18	3849	28	17	0	127	858	264	24
INVGOL.Z711INDDD3041.MACRESPOND	12	35	58	118	4	3	1027	66	22	492	23	13	0	53	171	253	11
SDTEST.Z711INDDD3041.MACRESPOND	12	30	10	318	23	22	769	15	7	1478	366	351	0	45	655	322	5
ORROCK.Z771RMMAT3042.MACRESPOND	11	22	59	4905	22	21	940	79	51	37486	198	214	0	57	1004	255	201
SDBPTE.Z728BPDKP3241.MACRESPOND	12	05	08	2581	101	100	851	18	50	18183	2949	852	2197	52	2203	187	207
CSCENE.Z713CTRJC4041.MACRESPOND	12	41	29	479	10	9	1015	35	35	4009	403	428	0	108	559	228	26
CSINVM.Z100INMJ66052.MACRESPOND	12	39	32	735	32	31	1029	43	6	6590	133	85	0	37	1072	248	16
MASTER.Z812SWIGK5071.MACRESPOND	12	36	22	660	42	42	1039	10	60	5991	18	3	0	7	291	124	59
MPLAN .Z791MLDCM6044.MACRESPOND	12	38	41	945	36	35	841	35	7	8290	429	402	0	36	849	117	0
CSCAP .Z724CXSRP3241.MACRESPOND	14	04	38	69	3	2	1002	63	6	1	21	13	0	71	156	267	2
PERT .Z752PP0156043.MACRESPOND	14	14	46	283	4	4	979	55	17	1000	49	13	0	93	341	267	8
SYSTEM.Z817SWKEN5071.MACRESPOND	14	11	58	634	7	6	1101	34	38	3000	41	2	0	46	259	222	17
SYSTEM.Z818SWKJF5071.MACRESPOND	14	18	20	253	3	2	1105	53	14	2000	34	4	0	72	130	256	3
CSINVM.Z100INBST6052.MACRESPOND	14	06	33	1024	31	30	1003	43	7	4733	178	134	0	36	944	231	19
ORORE .Z766GTRHW3042.MACRESPOND	14	20	34	344	5	4	1045	67	33	1000	24	56	0	62	198	239	21
SDMPAY.Z713MPRJC3041.MACRESPOND	14	10	02	1002	10	9	1003	41	40	4000	330	346	0	100	550	231	34
MPLAN .Z790MLWES6044.MACRESPOND	14	08	12	1347	75	74	938	12	7	4115	94	75	0	34	2450	67	12
CSBUDG.Z725BCIMP3041.MACRESPOND	14	24	39	453	16	15	902	19	11	1000	701	681	0	167	1878	110	7
CSBUDG.Z725BCCEL3041.MACRESPOND	14	33	17	444	3	3	994	74	9	2000	21	13	5	52	130	289	4
SDTEST.Z731STMCN3041.MACRESPOND	14	31	08	602	11	10	1010	57	9	2000	105	94	0	72	628	324	11
SYSTEM.Z818SWKJF5071.MACRESPOND	14	24	19	976	59	58	1047	10	40	3000	611	673	0	34	1273	127	54
CSBUDG.Z729BCDKP3241.MACRESPOND	14	11	53	1918	27	27	989	46	14	7000	260	108	235	70	1507	236	40
PERT .Z752PP0156043.MACRESPOND	14	41	34	195	4	3	1004	85	22	2000	25	17	0	84	263	244	14
SDTEST.Z731STSTC3041.MACRESPOND	14	42	11	276	6	6	950	64	17	2000	98	84	0	73	335	238	14
MASTER.Z812SWIGK5071.MACRESPOND	14	44	09	219	14	13	1016	16	37	2000	9	0	0	16	213	214	18
CCROCK.Z827RMFCN3042.MACRESPOND	14	11	01	2409	25	24	1107	47	246	10000	152	122	0	27	545	118	674
SYSTEM.Z817SWKEN5071.MACRESPOND	14	44	09	486	10	9	961	160	5	2000	619	604	0	111	355	393	12
ORROCK.Z827RMFCN3042.MACRESPOND	14	51	44	94	3	3	994	62	11	0	24	73	0	77	132	258	4
CSBUDG.Z725BCIMP3041.MACRESPOND	14	58	27	130	3	3	996	78	5	1000	26	14	0	57	131	244	2
SDMM .Z727MMB393241.MACRESPOND	14	14	41	2785	114	113	956	43	8	12000	3470	3263	0	89			
SYSTEM.Z817SWKEN5071.MACRESPOND	14	18	58	2695	66	65	1035	10	20	12000	1131	2478	1408	229			
SDTEST.Z731STSRR3041.MACRESPOND	14	57	50	424	7	6	1092	62	20	3000	80	55	0	70			
PERT .Z752PP0156043.MACRESPOND	15	03	28	96	4	3	976	73	8	1000	22	13	0	84			
SDMPAY.Z713MPRJC3041.MACRESPOND	14	52	44	863	35	34	936	34	6	4000	588	699	0	61			
SDTEST.Z731STSRR3041.MACRESPOND	15	06	53	203	6	5	1050	59	26	1000	38	20	0	55			
SDTEST.Z719STSTA3241.MACRESPOND	14	18	20	3125	32	31	1046	74	31	14000	266	114	0	72			
ORROCK.Z771RMMAT3042.MACRESPOND	14	27	29	2680	19	18	996	76	26	12000	188	206	0	74			
SDTEST.Z711STMCN3041.MACRESPOND	15	13	31	317	4	4	1062	79	18	1000	31	13	0	64			

MACRESPOND WORKLOAD GROUP

SUMMARY REPORT

USERNAME.JOBNAME.PROFILE	START TIME	ELAPSED TIME (SECS)	OCP SECS USER	OCP SECS TOTAL	M/C SPEED KIPS	VSI /SEC	USER PAGES /VSI	SYSTEM PAGE XFERS	FILE XFERS	USER DISC XFERS	USER TAPE XFERS	MM XFERS /SEC	DIREC TOR XFERS	OCC (KB)	NO OF RIROS
S DHH .Z727MHB393241.MACRESPOND	15 16 11	103	3	3	996	70	30	1000	22	14	0	53	122	258	12
PERT .Z752PP0156043.MACRESPOND	15 17 16	87	4	3	968	78	8	1000	22	13	0	88	238	262	4
ORORE .Z766GTRHW3042.MACRESPOND	15 21 02	68	4	3	1038	58	10	0	27	14	0	52	148	261	4
S DHPAY.Z713HPRJC3041.MACRESPOND	15 14 58	289	5	4	1007	77	31	1000	33	59	0	60	206	245	24
S DHPAY.Z713HPRJC3041.MACRESPOND	15 23 41	112	4	3	981	61	8	0	34	20	0	63	169	254	3
MASTER.Z812SWIGK5071.MACRESPOND	14 23 10	3925	143	142	803	7	42	16000	25358	19114	6208	185	1054	75	104
S DTEST.Z719STSTE3141.MACRESPOND	15 11 03	1197	73	73	941	16	19	4000	1953	1871	0	51	1852	182	52
ORORE .Z766GTRHW3042.MACRESPOND	15 32 19	116	4	3	1055	62	7	0	22	13	0	50	135	263	3
S UPPOR.Z822SWCSH5071.MACRESPOND	14 52 57	2659	48	47	1017	33	70	11000	134	110	0	52	2331	224	250
PERT .Z752PP0156043.MACRESPOND	15 39 30	108	4	3	982	81	7	1000	22	13	0	84	239	250	4
S DHPAY.Z713HPRJC3041.MACRESPOND	15 30 59	626	23	23	969	55	6	4000	689	874	0	87	1094	279	14
CSCAP .Z724CXSRP3241.MACRESPOND	15 08 28	2184	61	61	944	23	14	10000	1443	1475	0	58	2036	255	49
S DBPTE.Z728BPDKP3241.MACRESPOND	15 32 06	871	7	6	996	74	28	5000	36	14	5	62	363	223	30
ORCONV.Z764066DK3042.MACRESPOND	15 35 31	869	34	33	1028	38	3	6000	189	139	0	31	884	131	13
CCROCK.Z827RHFCN3042.MACRESPOND	14 55 26	3414	26	25	1197	10	86	18000	22	13	0	9	219	120	51
ENGINE.Z925HWHHH6769.MACRESPOND	15 32 18	1084	14	12	1501	46	13	6000	191	173	0	77	783	199	18
ORORE .Z766GTRHW3042.MACRESPOND	15 50 41	184	5	5	1085	54	16	1000	23	13	0	45	195	246	9
S DTEST.Z731STHCN3041.MACRESPOND	15 47 35	442	18	17	1030	19	19	4000	209	189	0	45	572	141	14
S DHH .Z715HHH393041.MACRESPOND	15 28 08	1615	32	32	1011	57	19	11000	482	1381	0	89	1431	202	77
S DHPAY.Z713HPRJC3041.MACRESPOND	15 42 43	917	111	110	964	10	4	8000	3467	3537	0	53	2326	210	10
PERT .Z752PP0156043.MACRESPOND	15 56 17	138	4	3	1003	66	9	1000	22	13	0	84	242	264	4
CSINV .Z723ISLAC3241.MACRESPOND	15 55 44	251	4	4	994	67	25	1000	94	76	0	73	206	239	15
S DHPAY.Z714HPCES3041.MACRESPOND	15 08 08	3125	11	11	972	40	152	17000	418	399	0	104	691	231	156
ORROCK.Z771RHHAT3042.MACRESPOND	15 31 17	1801	16	15	968	48	32	13000	74	123	0	45	560	154	54
S DHH .Z718HHI903041.MACRESPOND	15 51 46	700	8	7	1045	50	32	6000	56	31	0	60	400	238	27
S DISD .Z713SWRJC3241.MACRESPOND	15 58 35	338	6	5	999	62	41	2000	36	61	0	70	319	244	32
S DINV .Z723ISLAC3241.MACRESPOND	16 00 13	391	5	5	964	49	19	4000	26	14	0	151	710	233	10
CCROCK.Z827RHFCN3042.MACRESPOND	16 05 33	130	4	4	1055	80	8	2000	22	13	0	39	124	233	5
CSCAP .Z108CXFEV6080.MACRESPOND	15 58 35	1070	19	18	1000	46	13	8000	1051	1020	0	110	993	197	26
ORROCK.Z771RHHAT3042.MACRESPOND	16 10 38	543	6	6	1044	84	26	4000	37	18	0	46	251	240	29
S DTEST.Z731STSTC3041.MACRESPOND	16 04 41	959	24	23	1060	31	33	9000	178	182	0	42	812	167	55
ORROCK.Z771RHHAT3042.MACRESPOND	16 24 15	41	3	2	1017	66	4	0	46	13	0	58	108	333	1
CSINVP.Z100INBWT6052.MACRESPOND	16 04 06	274	17	16	1006	52	5	3000	100	73	0	37	531	240	8
ORROCK.Z771RHHAT3042.MACRESPOND	16 02 18	525	6	5	1037	79	24	4000	39	18	0	58	287	218	23
PERT .Z752PP0156043.MACRESPOND	16 11 27	91	4	3	1015	61	10	0	22	13	0	84	241	261	4
ORROCK.Z771RHHAT3042.MACRESPOND	16 11 56	51	3	2	1005	62	3	0	21	13	0	61	122	271	1
MPLAN .Z791MLDCH6044.MACRESPOND	14 05 43	8484	418	416	813	18	10	46743	5633	5382	0	27	5941	111	164
CSBUDG.Z725BCIHP3041.MACRESPOND	15 19 28	3484	113	112	938	38	5	23000	580	577	161	50	4857	219	43
S DISD .Z716SWROY3241.MACRESPOND	16 24 19	304	6	6	1107	68	28	2000	82	37	0	58	291	268	25
PERT .Z752PP0156043.MACRESPOND	16 20 31	94	4	3	979	70	9	1000	22	13	0	85	240	258	4
ORROCK.Z771RHHAT3042.MACRESPOND	16 21 29	93	3	2	1002	75	3	0	20	13	0	57	108	276	1
MASTER.Z812SWIGK5071.MACRESPOND	16 17 44	225	4	3	1165	35	33	3000	32	0	0	21	71	221	9
S DHPAY.Z714HPCES3041.MACRESPOND	16 16 44	702	5	5	1000	75	23	7000	57	37	0	65	268	236	18
ORROCK.Z771RHHAT3042.MACRESPOND	16 26 26	60	3	2	1018	91	3	1000	46	13	0	58	113	311	1
CSCENE.Z713CTRJC3404.MACRESPOND	12 40 40	14	2	0	0	0	1	0	0	0	0	0	0	0	0
S DHONE.Z723HHLAC6052.MACRESPOND	09 15 01	118	4	0	0	0	0	0	14	0	0	0	0	0	0
S DTEST.Z711INDDD3041.MACRESPOND	12 09 34	41	4	0	0	0	1	0	14	0	0	0	0	0	0

```
*******************************************************************************
       WORKLOAD GROUP         * START *ELAPSED* OCP SECS  * M/C * VSI *USER *SYSTEM* FILE * USER * USER *  MM  *DIREC * OCC *NO OF
   USERNAME.JOBNAME.PROFILE   * TIME  * TIME  ************SPEED* /SEC *PAGES* PAGE * XFERS* DISC * TAPE * XFERS*TOR  *(KB) *RIROS
                              *       *(SECS) *USER *TOTAL*KIPS *     */VSI *XFERS *      * XFERS* XFERS* /SEC *XFERS *
*******************************************************************************
```

THERE ARE 162 JOBS IN THIS WORKLOAD GROUP

WORKLOAD DEFINITION
****

PROFILE=MAC*
TIME=085800&125500&081400000&163000&10

SESSION END AT 01:05:54 ON 1981/06/05 ELAPSED TIME 58:24 SECONDS WITH A TOTAL OF 299 JOBS IN THIS SESSION

MACRESPOND WORKLOAD GROUP

SUMMARY REPORT

```
IITTTTTTRRRR
II  TT   RR RR              SESSION IDENTITY JUN041                    AT 08:58:41 ON 1981/06/04
II  TT   RRRR
II  TT   RR RR
```

NEW REPORT PROGRAM   -   DEVICE USAGE REPORT
****

```
**********************************************************************************************************************
* DEVICE   * VOLUME *   BLOCKSIZE    *   FILE   *  FILE TRANSFERS  * LOADER ACCESSES  *  OTHER ACCESSES  *
* ADDRESS  *  ID    *********************** ACCESSES ******************************************************************
*          *        * AVE   *  VAR   *          *  TOTAL  *  AS %  *  TOTAL  *  AS %  *  TOTAL  *  AS %  *
**********************************************************************************************************************
* 01010100  SYS003   2019   808739      174      8089     9.567       0      0.000      16      1.942*
* 01010000  SYS001   225    548351      281      3129     3.701     3875    90.920      34      4.126*
* 01040200  SYS002   2042  1501307       5      19225    22.737      173     4.059       6      0.728*
* 01040700  GFSA12   2048        0       25       4407     5.212       0      0.000       0      0.000*
* 01040400  GFSA02   1237  3757521      494     15189    17.964      103     2.417     305     37.015*
* 01010500  GFSA13   4977 23436852       74     13148    15.550       0      0.000      13      1.578*
* 01040600  GFSA03   1732   549202      313      8780    10.384      65      1.525      12      1.456*
* 01050001  DEV262   2048        0        1         3     0.004       0      0.000       0      0.000*
* 01040300  GFSA01   522    797491      128      3581     4.235      46      1.079     414     50.243*
* ********  DEV214   3422 17399215        6        136     0.161       0      0.000      20      2.427*
* 01050103  PRF008   6144        0        2       6179     7.308       0      0.000       0      0.000*
* 01050100  DEV262   7606 30904528        2       1120     1.325       0      0.000       0      0.000*
* 01040500  SYS004   2048        0       14        416     0.492       0      0.000       4      0.485*
* 01010700  GFSA11   2048        0        4         63     0.075       0      0.000       0      0.000*
* 01050102  DEV088  13165  *     0        1       1089     1.288       0      0.000       0      0.000*
**********************************************************************************************************************
* DEVICE   * VOLUME *   BLOCKSIZE    *   FILE   *  FILE TRANSFERS  * LOADER ACCESSES  *  OTHER ACCESSES  *
* ADDRESS  *  ID    *********************** ACCESSES ******************************************************************
*          *        * AVE   *  VAR   *          *  TOTAL  *  AS %  *  TOTAL  *  AS %  *  TOTAL  *  AS %  *
**********************************************************************************************************************
```

SESSION END AT 01:05:54 ON 1981/06/05 ELAPSED TIME 58024 SECONDS WITH A TOTAL OF 299 JOBS IN THIS SESSION

MACRESPOND WORKLOAD GROUP

DEVICE USAGE REPORT

INSTITUUT VIR TOEGEPASTE REKENAARWETENSKAP/INSTITUTE FOR APPLIED COMPUTER SCIENCE - UNIVERSITY OF STELLENBOSCH - SOUTH AFRICA

```
IITTTTTRRRR
II  TT  RR RR        SESSION IDENTITY JUN041            AT 08:58:41 ON 1981/06/04
II  TT  RRRR
II  TT  RR RR
```

STATISTICS FOR THIS WORKLOAD GROUP
****

		TOTAL	AVERAGE	VARIANCE
OCP TIME(SECS)	TOTAL	4487.235	27.70	2206.90
	USER	4618.781	28.51	2216.31
OCP UTILISATION			20.07	0.05
VIRTUAL STORE INTERRUPTS	TOTAL	139177	859.12	1241826.08
	USER	145533	898.35	1236576.03
INTER VSI TIME(MSECS)			32.24	2269.52
OCCUPANCY(KR)			211.78	4997.26
VIRTUAL STORE PAGES TRANSFERED	DISC TOTAL	3611081	22290.62	2517891892.50
	USER	3619424	22342.12	2528187777.65
	DRUM TOTAL	0	0.00	0.00
	USER	0	0.00	0.00
DIRECTOR TRANSFERS	LOCAL	13515	83.43	20276.07
	PUBLIC	97269	600.43	1145626.56
	LOWER	68911	425.38	2371414.25
	TOTAL	179695	1109.23	4203808.18
I/O-TRANSFERS	DISC	86307	532.76	3221214.20
	TAPE	24500	151.23	1260905.26
	FILES	89640	553.33	2379813684.61
MM-TRANSFERS	TOTAL	294024	1814.96	12531061.60
	USER	299091	1846.24	12539128.19
NO OF FILES ACCESSED		1987	12.27	308.43
NO OF RECORDS SPOOLED		0	0.00	0.00
NO OF RIROS		8622	53.22	14535.02
ELAPSED TIME		222620	1374.20	3078168.94

SESSION END AT 01:05:54 ON 1981/06/05 ELAPSED TIME 58024 SECONDS WITH A TOTAL OF 299 JOBS IN THIS SESSION

MACRESPOND WORKLOAD GROUP

STATISTICS REPORT

B-11

```
IITTTTTTRRRR
II  TT  RR RR        SESSION IDENTITY JUN041              AT 08:52:41 ON 1981/06/04
II  TT  RRRR
II  TT  RR RR
```

	TOTAL TIME(SECS)	NORMALISED
CONCURRENCY DISTRIBUTION		

LEVEL 0	4211	18.83
LEVEL 1	453	2.02
LEVEL 2	492	2.20
LEVEL 3	479	2.14
LEVEL 4	590	2.64
LEVEL 5	984	4.40
LEVEL 6	683	3.06
LEVEL 7	1295	5.79
LEVEL 8	2186	9.78
LEVEL 9	2361	10.56
LEVEL 10	1965	8.79
LEVEL 11	2403	10.75
LEVEL 12	1381	6.18
LEVEL 13	2074	9.27
LEVEL 14	4213	18.84
LEVEL 15	803	3.59
AVERAGE CONCURRENCY	9.96	

```
     -----------------------
SESSION END AT 01:05:54 ON 1981/06/05 ELAPSED TIME 58024 SECONDS WITH A TOTAL OF 299 JOBS IN THIS SESSION
```

MACRESPOND WORKLOAD GROUP

CONCURRENCY DISTRIBUTION REPORT

```
I IITTTTTTRRRR
II  TT  RR RR          SESSION IDENTITY JUN041              AT 08:58:41 ON 1981/06/04
II  TT  RRRR
II  TT  RR RR
```

```
                        NEW REPORT PROGRAM  -  SUMMARY REPORT
                        *********************************************
```

WORKLOAD GROUP USERNAME.JOBNAME.PROFILE	START TIME	ELAPSED TIME (SECS)	OCP SECS USER	TOTAL	M/C SPEED KIPS	VSI /VSI	USER /SEC PAGES	SYSTEM PAGE XFERS	FILE XFERS	USER DISC XFERS	USER TAPE XFERS	MM XFERS /SEC	DIREC TOR XFERS	OCC (KB)	NO OF RIROS
MASTER.OPERATOR .	08 58 50	62335	0	267	577	20	0	12A5	75	455	0	24	6009	258	0
MASTER.LP30ANYNS96' .SMCOS(LP30	19 46 19	49	0	0	1053	91	0	0	10	0	0	72	27	369	0
MASTER.LP10ANYNS64 .SMCOS(LP10	09 01 31	58036	424	503	877	37	5	146071	13094	12867	0	59	16739	93098	0
MASTER.LP30ANYNS96 .SMCOS(LP30	09 01 31	57655	368	408	881	46	6	146051	15819	15669	0	72	13601	107	0
MASTER.LP20ANYNS64 .SMCOS(LP20	09 01 51	57A69	292	398	916	36	5	145992	14092	13610	312	68	13036	107	0

WORKLOAD GROUP USERNAME.JOBNAME.PROFILE	START TIME	ELAPSED TIME (SECS)	OCP SECS USER	TOTAL	M/C SPEED KIPS	VSI /VSI	USER /SEC PAGES	SYSTEM PAGE XFERS	FILE XFERS	USER DISC XFERS	USER TAPE XFERS	MM XFERS /SEC	DIREC TOR XFERS	OCC (KB)	NO OF RIROS

```
                        THERE ARE   5 JOBS IN THIS WORKLOAD GROUP
```

```
WORKLOAD DEFINITION
*******************

PROFILE=SMCOS*
JOB=OPERATOR
```

```
        SESSION END AT 01:05:54 ON 1981/06/05 ELAPSEC TIME 58024 SECONDS WITH A TOTAL OF 299 JOBS IN THIS SESSION
```

SYSTEM WORKLOAD GROUP

SUMMARY REPORT

```
IITTTTTTRRRR
II  TT  RR RR         SESSION IDENTITY JUN041              AT 08:58:41 ON 1981/06/04
II  TT  RRRR
II  TT  RR RR
```

NEW REPORT PROGRAM - DEVICE USAGE REPORT
**************** *******************************

* DEVICE	* VOLUME *	BLOCKSIZE		* FILE	* FILE TRANSFERS		* LOADER ACCESSES		* OTHER ACCESSES	
* ADDRESS	* ID			ACCESSES						
*	*	* AVE	* VAR *	*	* TOTAL *	AS % *	TOTAL *	AS % *	TOTAL *	AS % *
* 01010100	SYS003	2048	0	277	14925	34.721	0	0.000	0	0.000*
* 01010000	SYS001	227	414504	0	0	0.000	27	90.000	0	0.000*
* 01040200	SYS002	2048	0	0	0	0.000	3	10.000	0	0.000*
* 01040700	GFSA12	2048	0	47	5241	12.193	0	0.000	0	0.000*
* 01040400	GFSA02	2048	0	126	8089	18.818	0	0.000	0	0.000*
* 01010500	GFSA13	2048	0	5	30	0.070	0	0.000	0	0.000*
* 01040600	GFSA03	2048	0	16	3786	8.808	0	0.000	0	0.000*
* 01040300	GFSA01	2048	0	6	151	0.351	0	0.000	0	0.000*
* 01050104	DEV317	2048	0	1	185	0.430	0	0.000	0	0.000*
* 01010600	GF0051	2048	0	37	6173	14.361	0	0.000	0	0.000*
* 01050100	DEV315	2048	0	1	121	0.281	0	0.000	0	0.000*
* 01040500	SYS004	2048	0	1	202	0.470	0	0.000	0	0.000*
* 01010700	GF0050	2048	0	13	4082	9.496	0	0.000	0	0.000*

* DEVICE	* VOLUME *	BLOCKSIZE		* FILE	* FILE TRANSFERS		* LOADER ACCESSES		* OTHER ACCESSES	
* ADDRESS	* ID			ACCESSES						
*	*	* AVE	* VAR *	*	* TOTAL *	AS % *	TOTAL *	AS % *	TOTAL *	AS % *

SESSION END AT 01:05:54 ON 1981/06/05 ELAPSED TIME 58024 SECONDS WITH A TOTAL OF 299 JOBS IN THIS SESSION

SYSTEM WORKLOAD GROUP

DEVICE USAGE REPORT

```
IITTTTTTRRRR
II  TT  RR RR          SESSION IDENTITY JUN041              AT 08:58:41 ON 1981/06/04
II  TT  RRRR
II  TT  RR RR
```

STATISTICS FOR THIS WORKLOAD GROUP
****

		TOTAL	AVERAGE	VARIANCE
OCP TIME(SECS)	TOTAL	1576.942	315.39	37356.88
	USER	1083.790	270.95	26667.52
OCP UTILISATION			2.72	0.00
VIRTUAL STORE INTERRUPTS	TOTAL	59527	12992.25	59122953.75
	USER	32225	8056.25	22838916.25
INTER VSI TIME(MSECS)			25.21	45.50
OCCUPANCY(KP)			23420.25	1618341055.75
VIRTUAL STORE PAGES TRANSFERED	DISC TOTAL	287707	71926.75	1866309615.75
	USER	69833	17458.25	1149585R2.25
	DRUM TOTAL	0	0.00	0.00
	USER	0	0.00	0.00
DIRECTOR TRANSFERS	LOCAL	7515	1878.75	4215587.25
	PUBLIC	26281	6570.25	14620134.75
	LOWER	9607	2401.75	2785719.75
	TOTAL	43403	10850.75	41041061.25
I/O-TRANSFERS	DISC	42126	10531.50	38008255.25
	TAPE	312	78.00	18252.00
	FILES	43015	10753.75	598993128.50
PP-TRANSFERS	TOTAL	85841	21460.25	154160727.25
	USER	68893	17223.25	102780842.75
NO OF FILES ACCESSED		536	134.00	6322.50
NO OF RECORDS SPOOLED		711578	177894.50	11236756912.25
NO OF RIROS		0	0.00	0.00
ELAPSED TIME		235944	43402.25	626519665.75

SESSION END AT 01:05:54 ON 1981/06/05 ELAPSED TIME 58024 SECONDS WITH A TOTAL OF 299 JOBS IN THIS SESSION

SYSTEM WORKLOAD GROUP

STATISTICS REPORT

INSTITUUT VIR TOEGEPASTE REKENAARWETENSKAP/INSTITUTE FOR APPLIED COMPUTER SCIENCE - UNIVERSITY OF STELLENBOSCH - SOUTH AFRICA

```
IITTTTTTRRRR
II   TT   RR RR        SESSION IDENTITY JUN041              AT 09:59:41 ON 1981/06/04
II   TT   RRRR
II   TT   RR RR
```

```
                                           TOTAL TIME(SECS)   NORMALISED
                    CONCURRENCY DISTRIBUTION
                    --------------------------
                              LEVEL  0              0          0.00
                              LEVEL  1           4299          6.90
                              LEVEL  2            147          0.24
                              LEVEL  3            254          0.40
                              LEVEL  4          57586         92.38
                              LEVEL  5             49          0.08
                    AVERAGE CONCURRENCY            3.79
                    -------------------------
```

SESSION END AT 01:05:54 ON 1981/06/05 ELAPSED TIME 58024 SECONDS WITH A TOTAL OF 299 JOBS IN THIS SESSION

B-16

SYSTEM WORKLOAD GROUP

CONCURRENCY DISTRIBUTION REPORT

APPENDIX C - GFSA JOBNAME STANDARDS

Jobs run on the GFSA computer are classified in two ways - firstly by a user name (up to ten characters in length), and secondly by a fixed format 13 character job name. The format is as follows:

    :username.MAAASSXXXJJCC

The meanings of the various fields are as follows:

    M   = mode (Z = MACRESPOND (i.e. terminal);

                    B = BATCH (job inserted via cards);

                    R = REMOTE BATCH (job inserted via a terminal))

    AAA = 3 digit account code (one code per user)

    SS   = system code (eg. ST : stores;

                       FB : financial budgeting;

                       SW : software)

    XXX = any 3 character job identifier

    JJ   = job code (eg. compile, test, edit)

    CC   = cost centre (for accounting use)

APPENDIX D - BACKGROUND PROFILES

The facility exists under VME/B for each site to define its own background profiles. This allows the system manager to control his resources in an optimal fashion by setting up specific profiles to do specific tasks, eg. preloading the compilers into a profile designed to run compilations.

The following profiles have been defined at GFSA:

   BATCH - a general profile for work that does not fall into any of
            the other categories;

   COMP - a profile to be used for compilations. The compilers are
            preloaded into the profile;

   DEVEL - used for development work for quick testing to avoid the
            need to run programs from the terminals;

   MILL - a profile placed at the bottom of the OCP priority queue
            for jobs that use a lot of OCP (i.e. mill) time;

   BMILL - similar to the mill profile, but with a larger main store
            allocation - designed for mill intensive jobs that also
            require a lot of main store.

APPENDIX E - EXAMPLE OF SNAP OUTPUT FORMAT

It is impossible to reproduce all the SNAP listings generated during
the calibration and validation stages. This appendix contains one
such listing, namely that produced during the final calibration run
with concurrencies of 4 and 3 for the BATCH and SYSTEM workloads
respectively.

The output provides details of the network parameters together with
a list of the chains identified by SNAP. The analysis is then
presented for each class in the model as well as an analysis
independent of class. Finally the response time for each open chain
is given.

S T O C H A S T I C   N E T W O R K   A N A L Y S I S   P A C K A G E

DEVELOPED BY THE PERFORMANCE EVALUATION RESEARCH GROUP
DEPARTMENT OF COMPUTER SCIENCE
UNIVERSITY OF STELLENBOSCH
STELLENBOSCH 7600
SOUTH AFRICA

COPYRIGHT (C) 1977 BY PETER TEUNISSEN

N E T W O R K   D E S C R I P T I O N

NUMBER OF SERVICE CENTERS 26

NUMBER OF CUSTOMER CLASSES = 5

NAMES OF SERVICE CENTERS:

1=OCP	2=SCHED	3=DC10	4=DC40	5=SYS002	6=SYS003
7=SYS001	8=GFSA13	9=ED16	10=ED17	11=GFSA01	12=GFSA02
13=GFSA03	14=GFSA12	15=SYS004	16=MC00	17=MC10	18=LP10
19=LP20	20=LP30	21=UTERM	22=UVSI	23=UDISC	24=URIRO
25=UTAPE	26=UPAPER				

TYPE OF EACH CENTER:

CENTER 1= 2	CENTER 2= 3	CENTER 3= 2	CENTER 4= 2	CENTER 5= 1
CENTER 6= 1	CENTER 7= 1	CENTER 8= 1	CENTER 9= 1	CENTER 10= 1
CENTER 11= 1	CENTER 12= 1	CENTER 13= 1	CENTER 14= 1	CENTER 15= 1
CENTER 16= 1	CENTER 17= 1	CENTER 18= 1	CENTER 19= 1	CENTER 20= 1
CENTER 21= 3	CENTER 22= 3	CENTER 23= 3	CENTER 24= 3	CENTER 25= 3
CENTER 26= 3				

SERVICE RATES:

	1	2	3	4	5
OCP	91.863	91.863	83.448	83.448	543.770
SCHED	2.000	2.000	0.004	0.004	0.100
DC1C	1.866	403.000	403.000	1.399	403.000
DC4C	1.866	403.000	403.000	1.399	403.000
SYSC02	26.000	26.000	26.000	26.C00	26.000
SYSC03	26.000	26.000	26.000	26.000	26.000
SYSC01	26.000	26.000	26.000	26.000	26.000
GFSA13	26.000	26.000	26.000	26.000	26.000
ED16	26.000	26.000	26.000	26.000	26.000
ED17	26.000	26.000	26.000	26.C00	26.000
GFSA01	26.000	26.000	26.000	26.000	26.000
GFSA02	26.000	26.000	26.000	26.000	26.000
GFSA03	26.000	26.000	26.000	26.000	26.000
GFSA12	26.000	26.000	26.000	26.000	26.000
SYSC04	26.000	26.000	26.000	26.000	26.000

SERVICE RATES:

	1	2	3	4	5
MCOC	67.000	67.000	67.000	67.000	67.000
MC1C	67.000	67.000	67.000	67.000	67.000
LP1C	6.400	6.400	6.400	6.400	6.400
LP2C	6.400	6.400	6.400	6.400	6.400
LP3C	6.400	6.400	6.400	6.400	6.400
UTERM	9999.000	9999.000	9999.000	9999.000	9999.000
UVSI	9999.000	9999.000	9999.000	9999.000	9999.000
UDISC	9999.000	9999.000	9999.000	9999.000	9999.000
URIRO	9999.000	9999.000	9999.000	9999.000	9999.000
UTAPE	9999.000	9999.000	9999.000	9999.000	9999.000
UPAPER	9999.000	9999.000	9999.000	9999.000	9999.000

EXTERNAL ARRIVAL RATES TO CENTERS:

	1	2	3	4	5
OCP	0.386	0.000	0.000	0.000	0.000
SCHED	0.000	0.000	0.000	0.000	0.000
DC1C	0.000	0.000	0.000	0.000	0.000
DC4C	0.000	0.000	0.000	0.000	0.000
SYSCO2	0.000	0.000	0.000	0.000	0.000
SYSCO3	0.000	0.000	0.000	0.000	0.000
SYSCO1	0.000	0.000	0.000	0.000	0.000
GFSA13	0.000	0.000	0.000	0.000	0.000
ED16	0.000	0.000	0.000	0.000	0.000
ED17	0.000	0.000	0.000	0.000	0.000
GFSA01	0.000	0.000	0.000	0.000	0.000
GFSA02	0.000	0.000	0.000	0.000	0.000
GFSA03	0.000	0.000	0.000	0.000	0.000
GFSA12	0.000	0.000	0.000	0.000	0.000
SYSCO4	0.000	0.000	0.000	0.000	0.000

EXTERNAL ARRIVAL RATES TO CENTERS:

	1	2	3	4	5
MCOC	0.000	0.000	0.000	0.000	0.000
MC1C	0.000	0.000	0.000	0.000	0.000
LP1C	0.000	0.000	0.000	0.000	0.000
LP2C	0.000	0.000	0.000	0.000	0.000
LP3C	0.000	0.000	0.000	0.000	0.000
UTERM	0.000	0.000	0.000	0.000	0.000
UVSI	0.000	0.000	0.000	0.000	0.000
UDISC	0.000	0.000	0.000	0.000	0.000
URIRO	0.000	0.000	0.000	0.000	0.000
UTAPE	0.000	0.000	0.000	0.000	0.000
UPAPER	0.000	0.000	0.000	0.000	0.000

BRANCHING PROBABILITIES

1	2	21	2	0.0004000	1	2	22	2	0.3376000	1	2	23	2	0.6207000
1	2	25	2	0.0204000	21	2	2	2	1.0000000	2	2	1	2	1.0000000
22	2	5	2	0.5000000	22	2	6	2	0.5000000	23	2	5	2	0.1038000
23	2	6	2	0.0317000	23	2	7	2	0.6509000	23	2	8	2	0.0514500
23	2	10	2	0.0002500	23	2	11	2	0.0216000	23	2	12	2	0.0764000
23	2	13	2	0.0451000	23	2	14	2	0.0172000	23	2	15	2	0.0016000
25	2	16	2	0.5000000	25	2	17	2	0.5000000	5	2	4	2	1.0000000
6	2	3	2	1.0000000	7	2	3	2	1.0000000	8	2	3	2	1.0000000
10	2	3	2	1.0000000	11	2	4	2	1.0000000	12	2	4	2	1.0000000
13	2	4	2	1.0000000	14	2	4	2	1.0000000	15	2	4	2	1.0000000
3	2	1	2	1.0000000	4	2	1	2	1.0000000	16	2	1	2	1.0000000
17	2	1	2	1.0000000	1	1	2	1	1.0000000	2	1	24	1	1.0000000
24	1	15	1	1.0000000	15	1	4	1	1.0000000	4	1	1	2	1.0000000
1	3	21	3	0.0002200	1	3	22	3	0.1349000	1	3	23	3	0.7401200
1	3	24	3	0.0004600	1	3	25	3	0.1243000	21	3	2	3	1.0000000
22	3	5	3	0.5000000	22	3	6	3	0.5000000	23	3	5	3	0.0622000
23	3	6	3	0.1085000	23	3	7	3	0.4618000	23	3	8	3	0.0251000
23	3	11	3	0.0276000	23	3	12	3	0.1047000	23	3	13	3	0.0063000
23	3	14	3	0.1959000	23	3	15	3	0.0079000	24	3	15	4	1.0000000
25	3	16	3	0.5000000	25	3	17	3	0.5000000	5	3	4	3	1.0000000
6	3	3	3	1.0000000	7	3	3	3	1.0000000	8	3	3	3	1.0000000
11	3	4	3	1.0000000	12	3	4	3	1.0000000	13	3	4	3	1.0000000
14	3	4	3	1.0000000	15	3	4	3	1.0000000	15	4	4	4	1.0000000
4	3	1	3	1.0000000	3	3	1	3	1.0000000	4	4	2	3	1.0000000

16	3	1	3	1.0000000	17	3	1	3	1.0000000	2	3	1	3	1.0000000
1	5	21	5	0.0000060	1	5	22	5	0.0694000	1	5	23	5	0.1004000
1	5	25	5	0.0003570	1	5	26	5	0.8298370	21	5	2	5	1.0000000
22	5	5	5	0.5000000	22	5	6	5	0.5000000	23	5	5	5	0.0504000
23	5	6	5	0.1734000	23	5	7	5	0.4537000	23	5	8	5	0.0004000
23	5	9	5	0.0717000	23	5	10	5	0.0474000	23	5	11	5	0.0017000
23	5	12	5	0.0940000	23	5	13	5	0.0440000	23	5	14	5	0.0609000
23	5	15	5	0.0024000	25	5	16	5	0.5000000	25	5	17	5	0.5000000
26	5	18	5	0.3300000	26	5	19	5	0.3300000	26	5	20	5	0.3400000
5	5	4	5	1.0000000	6	5	3	5	1.0000000	7	5	3	5	1.0000000
8	5	3	5	1.0000000	9	5	3	5	1.0000000	10	5	3	5	1.0000000
11	5	4	5	1.0000000	12	5	4	5	1.0000000	13	5	4	5	1.0000000
14	5	4	5	1.0000000	15	5	4	5	1.0000000	16	5	1	5	1.0000000
17	5	1	5	1.0000000	18	5	1	5	1.0000000	19	5	1	5	1.0000000
20	5	1	5	1.0000000	3	5	1	5	1.0000000	4	5	1	5	1.0000000
2	5	1	5	1.0000000										

DEPARTURE PROBABILITIES

	1	2	3	4	5
OCP	0.000	0.021	0.000	1.000	0.000
SCHED	0.000	0.000	0.000	1.000	0.000
DC1C	1.000	0.000	0.000	1.000	0.000
DC4C	0.000	0.000	0.000	0.000	0.000
SYSC02	1.000	0.000	0.000	1.000	0.000
SYSC03	1.000	0.000	0.000	1.000	0.000
SYSC01	1.000	0.000	0.000	1.000	0.000
GFSA13	1.000	0.000	0.000	1.000	0.000
ED16	1.000	1.000	1.000	1.000	0.000
ED17	1.000	0.000	1.000	1.000	0.000
GFSA01	1.000	0.000	0.000	1.000	0.000
GFSA02	1.000	0.000	0.000	1.000	0.000
GFSA03	1.000	0.000	0.000	1.000	0.000
GFSA12	1.000	0.000	0.000	1.000	0.000
SYSC04	0.000	0.000	0.000	0.000	0.000

E-9

DEPARTURE PROBABILITIES

	1	2	3	4	5
MC00	1.000	0.000	0.000	1.000	0.000
MC1C	1.000	0.000	0.000	1.000	0.000
LP1C	1.000	1.000	1.000	1.000	0.000
LP2C	1.000	1.000	1.000	1.000	0.000
LP3C	1.000	1.000	1.000	1.000	0.000
UTERM	1.000	0.000	0.000	1.000	0.000
UVSI	1.000	0.000	0.000	1.000	0.000
UDISC°	1.000	-0.000	-0.000	1.000	-0.000
URIRO	0.000	1.000	0.000	1.000	1.000
UTAPE	1.000	0.000	0.000	1.000	0.000
UPAPER	1.000	1.000	1.000	1.000	0.000

THE FOLLOWING SUBCHAINS HAVE BEEN IDENTIFIED

NO OF SUBCHAINS=   3

*CHAIN NO 1 (CLOSED)

         (1, 3)   (21, 3)   (22, 3)   (23, 3)   (24, 3)   (25, 3)   (2, 3)   (5, 3)   (6, 3)   (7, 3)
         (8, 3)   (11, 3)   (12, 3)   (13, 3)   (14, 3)   (15, 3)   (15, 4)   (16, 3)   (17, 3)   (4, 3)
         (3, 3)   (4, 4)

*CHAIN NO 2 (CLOSED)

         (1, 5)   (21, 5)   (22, 5)   (23, 5)   (25, 5)   (26, 5)   (2, 5)   (5, 5)   (6, 5)   (7, 5)
         (8, 5)   (9, 5)   (10, 5)   (11, 5)   (12, 5)   (13, 5)   (14, 5)   (15, 5)   (16, 5)   (17, 5)
         (18, 5)   (19, 5)   (20, 5)   (4, 5)   (3, 5)

*CHAIN NO 3 (OPEN)

         (1, 1)   (2, 1)   (24, 1)   (15, 1)   (4, 1)   (1, 2)   (21, 2)   (22, 2)   (23, 2)   (25, 2)
         (2, 2)   (5, 2)   (6, 2)   (7, 2)   (8, 2)   (10, 2)   (11, 2)   (12, 2)   (13, 2)   (14, 2)
         (15, 2)   (16, 2)   (17, 2)   (4, 2)   (3, 2)

NUMBER OF CUSTOMERS IN CLOSED SUBCHAINS:

             4 IN CHAIN 1   3 IN CHAIN 2
NCONST COMPLETED. NORMALIZING CONSTANT=   0.588454E-06
COMPLETED CALCULATIONS FOR CENTER   26
NCONST COMPLETED. NORMALIZING CONSTANT=   0.588313E-06
COMPLETED CALCULATIONS FOR CENTER   25
NCONST COMPLETED. NORMALIZING CONSTANT=   0.588431E-06
COMPLETED CALCULATIONS FOR CENTER   24
NCONST COMPLETED. NORMALIZING CONSTANT=   0.586997E-06
COMPLETED CALCULATIONS FOR CENTER   23
NCONST COMPLETED. NORMALIZING CONSTANT=   0.587904E-06
COMPLETED CALCULATIONS FOR CENTER   22
NCONST COMPLETED. NORMALIZING CONSTANT=   0.588454E-06
COMPLETED CALCULATIONS FOR CENTER   21
NCONST COMPLETED. NORMALIZING CONSTANT=   0.239696E-06
COMPLETED CALCULATIONS FOR CENTER   20
NCONST COMPLETED. NORMALIZING CONSTANT=   0.249954E-06
COMPLETED CALCULATIONS FOR CENTER   19
NCONST COMPLETED. NORMALIZING CONSTANT=   0.249954E-06

```
COMPLETED CALCULATIONS FOR CENTER    18
NCONST COMPLETED. NORMALIZING CONSTANT=    0.577937E-06
COMPLETED CALCULATIONS FOR CENTER    17
NCONST COMPLETED. NORMALIZING CONSTANT=    0.577937E-06
COMPLETED CALCULATIONS FOR CENTER    16
NCONST COMPLETED. NORMALIZING CONSTANT=    0.576924E-06
COMPLETED CALCULATIONS FOR CENTER    15
NCONST COMPLETED. NORMALIZING CONSTANT=    0.528971E-06
COMPLETED CALCULATIONS FOR CENTER    14
NCONST COMPLETED. NORMALIZING CONSTANT=    0.573708E-06
COMPLETED CALCULATIONS FOR CENTER    13
NCONST COMPLETED. NORMALIZING CONSTANT=    0.537362E-06
COMPLETED CALCULATIONS FOR CENTER    12
NCONST COMPLETED. NORMALIZING CONSTANT=    0.575312E-06
COMPLETED CALCULATIONS FOR CENTER    11
NCONST COMPLETED. NORMALIZING CONSTANT=    0.586941E-06
COMPLETED CALCULATIONS FOR CENTER    10
NCONST COMPLETED. NORMALIZING CONSTANT=    0.586263E-06
COMPLETED CALCULATIONS FOR CENTER     9
NCONST COMPLETED. NORMALIZING CONSTANT=    0.568292E-06
COMPLETED CALCULATIONS FOR CENTER     8
NCONST COMPLETED. NORMALIZING CONSTANT=    0.280523E-06
COMPLETED CALCULATIONS FOR CENTER     7
NCONST COMPLETED. NORMALIZING CONSTANT=    0.439699E-06
COMPLETED CALCULATIONS FOR CENTER     6
NCONST COMPLETED. NORMALIZING CONSTANT=    0.437330E-06
COMPLETED CALCULATIONS FOR CENTER     5
NCONST COMPLETED. NORMALIZING CONSTANT=    0.444814E-06
COMPLETED CALCULATIONS FOR CENTER     4
NCONST COMPLETED. NORMALIZING CONSTANT=    0.557451E-06
COMPLETED CALCULATIONS FOR CENTER     3
NCONST COMPLETED. NORMALIZING CONSTANT=    0.573468E-08
COMPLETED CALCULATIONS FOR CENTER     2
NCONST COMPLETED. NORMALIZING CONSTANT=    0.338974E-06
COMPLETED CALCULATIONS FOR CENTER     1


NORMALIZING CONSTANT=0.588454E-06
```

N E T W O R K   S T A T I S T I C S

ANALYSIS FOR CUSTOMERS OF CLASS   1

CENTER	UTILIZATION	E(QUEUE)	E(TIME SPENT)	CYCLE TIME	RATE
1 OCP	0.00419864	0.0071	0.0185	999999900000000.0000	0.3857
2 SCHED	9.09090900	0.1929	0.5000	999999900000000.0000	0.3857
4 DC40	0.20673200	0.2734	0.7090	999999900000000.0000	0.3857
15 SYS004	0.01483461	0.0151	0.0392	999999900000000.0000	0.3857
24 URIRO	9.09090900	0.0000	0.0001	999999900000000.0000	0.3857

ANALYSIS FOR CUSTOMERS OF CLASS 2

CENTER	UTILIZATION	E(QUEUE)	E(TIME SPENT)	CYCLE TIME	RATE
1 OCP	0.20088930	0.3409	0.0185	999999900000000.0000	18.4543
2 SCHED	9.09090900	0.0037	0.5000	999999900000000.0000	0.0074
3 DC10	0.02860098	0.0302	0.0026	999999900000000.0000	11.5262
4 DC40	0.01528181	0.0202	0.0033	999999900000000.0000	6.1586
5 SYS002	0.16554120	0.2223	0.0516	999999900000000.0000	4.3041
6 SYS003	0.13377680	0.1784	0.0513	999999900000000.0000	3.4782
7 SYS001	0.28676120	0.5794	0.0777	999999900000000.0000	7.4558
8 GFSA13	0.02266686	0.0235	0.0398	999999900000000.0000	0.5893
10 ED17	0.00011014	0.0001	0.0386	999999900000000.0000	0.0029
11 GFSA01	0.00951612	0.0097	0.0393	999999900000000.0000	0.2474
12 GFSA02	0.03365883	0.0368	0.0421	999999900000000.0000	0.8751
13 GFSA03	0.01986930	0.0204	0.0394	999999900000000.0000	0.5166
14 GFSA12	0.00757764	0.0084	0.0427	999999900000000.0000	0.1970
15 SYS004	0.00070490	0.0007	0.0392	999999900000000.0000	0.0183
16 MC00	0.00280946	0.0029	0.0152	999999900000000.0000	0.1882
17 MC10	0.00280946	0.0029	0.0152	999999900000000.0000	0.1882
21 UTERM	9.09090900	0.0000	0.0001	999999900000000.0000	0.0074
22 UVSI	9.09090900	0.0006	0.0001	999999900000000.0000	6.2302
23 UDISC	9.09090900	0.0011	0.0001	999999900000000.0000	11.4546
25 UTAPE	9.09090900	0.0000	0.0001	999999900000000.0000	0.3765

ANALYSIS FOR CUSTOMERS OF CLASS 3

CENTER	UTILIZATION	E(QUEUE)	E(TIME SPENT)	CYCLE TIME	RATE
1 OCP	0.19414340	0.3071	0.0190	0.2279	16.2009
2 SCHED	9.09090900	2.8248	256.4102	106.6748	0.0110
3 DC10	0.02042677	0.0215	0.0026	0.4833	8.2320
4 DC40	0.01474980	0.0194	0.0033	0.6697	5.9442
5 SYS002	0.07071376	0.0929	0.0505	2.1251	1.8386
6 SYS003	0.09206659	0.1194	0.0499	1.6212	2.3937
7 SYS001	0.21297130	0.3938	0.0711	0.6513	5.5373
8 GFSA13	0.01157562	0.0120	0.0397	13.2508	0.3010
11 GFSA01	0.01272854	0.0130	0.0392	12.0475	0.3309
12 GFSA02	0.04828557	0.0522	0.0416	3.1446	1.2554
13 GFSA03	0.00290543	0.0030	0.0394	52.9119	0.0755
14 GFSA12	0.09034485	0.0980	0.0417	1.6612	2.3490
15 SYS004	0.00364331	0.0037	0.0392	42.1878	0.0947
16 MC00	0.01502823	0.0152	0.0151	3.9575	1.0069
17 MC10	0.01502823	0.0152	0.0151	3.9575	1.0069
21 UTERM	9.09090900	0.0000	0.0001	1122.2650	0.0036
22 UVSI	9.09090900	0.0002	0.0001	1.8301	2.1855
23 UDISC	9.09090900	0.0012	0.0001	0.3335	11.9907
24 URIRO	9.09090900	0.0000	0.0001	536.7351	0.0075
25 UTAPE	9.09090900	0.0002	0.0001	1.9862	2.0138

ANALYSIS FOR CUSTOMERS OF CLASS 4

CENTER	UTILIZATION	E(QUEUE)	E(TIME SPENT)	CYCLE TIME	RATE
4 DC40	0.00532583	0.0070	0.9393	535.7981	0.0075
15 SYS004	0.00028663	0.0003	0.0392	536.6985	0.0075

ANALYSIS FOR CUSTOMERS OF CLASS 5

CENTER	UTILIZATION	E(QUEUE)	E(TIME SPENT)	CYCLE TIME	RATE
1 OCP	0.02472300	0.0416	0.0031	0.2201	13.4436
2 SCHED	9.09090900	0.0008	10.0000	37181.9700	0.0001
3 DC10	0.00365810	0.0039	0.0026	2.0324	1.4742
4 DC40	0.00200625	0.0027	0.0033	3.7072	0.8085
5 SYS002	0.02055855	0.0275	0.0514	5.5611	0.5345
6 SYS003	0.02694386	0.0357	0.0510	4.2314	0.7005
7 SYS001	0.02355299	0.0472	0.0771	4.8218	0.6124
8 GFSA13	0.00002077	0.0000	0.0398	5556.5630	0.0005
9 ED16	0.00372219	0.0037	0.0386	30.9606	0.0968
10 ED17	0.00246069	0.0025	0.0385	46.8526	0.0640
11 GFSA01	0.00008825	0.0001	0.0393	1307.3960	0.0023
12 GFSA02	0.00487984	0.0053	0.0421	23.6031	0.1269
13 GFSA03	0.00228419	0.0023	0.0394	50.4751	0.0594
14 GFSA12	0.00316153	0.0035	0.0427	36.4538	0.0822
15 SYS004	0.00012459	0.0001	0.0392	926.0591	0.0032
16 MCO0	0.00003582	0.0000	0.0152	1250.1420	0.0024
17 MC10	0.00003582	0.0000	0.0152	1250.1420	0.0024
18 LP10	0.57523540	0.9267	0.2517	0.5632	3.6815
19 LP20	0.57523540	0.9267	0.2517	0.5632	3.6815
20 LP30	0.59266670	0.9683	0.2553	0.5356	3.7931
21 UTERM	9.09090900	0.0000	0.0001	37192.1300	0.0001
22 UVSI	9.09090900	0.0001	0.0001	3.2154	0.9330
23 UDISC	9.09090900	0.0001	0.0001	2.2225	1.3497
25 UTAPE	9.09090900	0.0000	0.0001	625.0774	0.0048
26 UPAPER	0.00037191	0.0011	0.0001	0.2688	11.1561

ANALYSIS INDEPENDENT OF CLASS

CENTER	UTILIZATION	E(QUEUE)	E(TIME SPENT)	CYCLE TIME	RATE	
1 OCP	0.42395440	0.6967	0.0144	9999998000000000.0000		48.4845
2 SCHED	9.09090900	3.0221	7.4772	9999998000000000.0000		0.4042
3 DC10	0.05268584	0.0555	0.0026	9999998000000000.0000		21.2324
4 DC40	0.24409560	0.3227	0.0243	9999998000000000.0000		13.3044
5 SYS002	0.25681350	0.3426	0.0513	9999998000000000.0000		6.6772
6 SYS003	0.25278720	0.3335	0.0507	9999998000000000.0000		6.5725
7 SYS001	0.52328540	1.0204	0.0750	9999998000000000.0000		13.6054
8 GFSA13	0.03426325	0.0354	0.0398	9999998000000000.0000		0.8908
9 ED16	0.00372219	0.0037	0.0386	30.9606	0.0968	0.0668
10 ED17	0.00257083	0.0026	0.0385	9999998000000000.0000		0.5807
11 GFSA01	0.02233291	0.0228	0.0393	9999998000000000.0000		2.2574
12 GFSA02	0.08682418	0.0944	0.0418	9999998000000000.0000		0.6515
13 GFSA03	0.02505891	0.0257	0.0394	9999998000000000.0000		2.6282
14 GFSA12	0.10108390	0.1099	0.0418	9999998000000000.0000		0.5094
15 SYS004	0.01959404	0.0200	0.0392	9999998000000000.0000		1.1975
16 MC00	0.01787350	0.0181	0.0151	9999998000000000.00C0		1.1975
17 MC10	0.01787350	0.0181	0.0151	9999998000000000.0000		
18 LP10	0.57523540	0.9267	0.2517	0.5632	3.6815	
19 LP20	0.57523540	0.9267	0.2517	0.5632	3.6815	
20 LP30	0.59266670	0.9683	0.2553	0.5356	3.7931	
21 UTERM	9.09090900	0.0000	0.0001	9999998000000000.00C0		0.0110
22 UVSI	9.09090900	0.0009	0.0001	9999998000000C0CCC.00C0		9.3487
23 UDISC	9.09090900	0.0025	0.0001	9999998000000000C.0GC0		24.7950
24 URIRO	9.09090900	0.0000	0.0001	9999998000000000.0000		0.3932
25 UTAPE	9.09090900	0.0002	0.0001	9999998000000000.0000		2.3951
26 UPAPER	0.00037191	0.0011	0.0001	0.2688	11.1561	

EXPECTED TIME SPENT IN SYSTEM FOR OPEN SUBCHAINS

CHAIN NO   E(TIME SPENT)

3   5.1094

AVERAGE   5.1094

REFERENCES

FERRARI D   Computer Systems Performance Evaluation (1978)

ITR Publications:
      An introduction to computer performance modelling (1980)
      An introduction to system modelling using the Stochastic
      Network Analysis Package (1977)
      Stochastic Network Analysis Program User's Manual (1977)

Numerous ICL manuals

Project documentation