



THE UNIVERSITY OF QUEENSLAND

THE PRENTICE CENTRE

18 March 1993

**Capacity Planning
for
Administrative Computing:
Recommendations for 1993**

Introduction and Executive Summary

This report is the outcome of deliberations of a committee composed of staff drawn from The Prentice Centre and the Administrative Information Systems section. The purpose of the report is to detail the capital expenditure items that will be necessary to ensure adequate performance and operational security for administrative computer systems in 1993.

This report does not consider requirements for 1994 and beyond. That is properly the object of a more detailed study which examines issues of processor, I/O, and communications architecture, software platforms, CSMAC compliance, environmental security, and other topics.

The Committee recommends: (in order of urgency)

Purchase 256MB of memory for UQADM	\$148 000
Purchase 5 Micro Tech MDI230 disks for UQADM (or equivalent)	\$58 000
Purchase 2 second-hand LP29 printers for UQADM	\$16 000
OR	
Make more use of laser printing facilities. (per annum)	\$10 000
Install TA867 42 GByte cartridge loader tape subsystem	\$32 500
OR	
Upgrade 2.3GB tape Unit to 5GB unit	\$16 000
Purchase disk drive for UQEVE development database systems	\$5 000
Remove UQEVE from VAXCluster, operate standalone	\$10 000
Relocate B1000 lineprinters to JD Story building	\$4 500
Purchase 128MB of memory for UQVAX	\$74 000
Upgrade HSC1 from an HSC50 to an HSC95	\$39 500

(Note that as at 16-Mar-93, the first item on the list had been fully implemented, and the necessary equipment for item 2 had all either been installed or was on order.)

Methodology

The first phase of the task involved a survey of current disk use and current system loads and projections estimating increases in these over 1993. AIS provided current figures along with estimates of:

- disk space requirements growth in administrative applications in 1993.
- growth in number of users of administrative applications in 1993.

Using these estimates, simulation software and disk capacity calculations were used to:

- estimate amount of extra disk and memory required for adequate performance under *current* load conditions;
- apply growth estimates supplied by AIS to arrive at projections for memory, CPU, and mass storage requirements for the end of 1993;
- list viable options with cost estimates and cost/benefit analysis.

Capacity planning for the upgrade of *existing* computer systems is performed in three stages:

Memory: Estimate the amount of real memory needed to provide adequate performance given the number of users and expected size of images they will be running.

Filestore: Estimate the total required pagefile and swapfile space. Allow 4-500 000 blocks of free space per spindle. Estimate growth of databases. Estimate extra storage quota required due to new users. Allow for a reasonable amount of reorganisation and scratch space to manage the databases.

Processor: Using estimates already arrived at as a base, ascertain whether or not a CPU upgrade will be necessary.

In addition, capacity requirements for **printing** and **backup** were considered.

Backup Processor: Specify what upgrade, if any, will need to be made to UQVAX in order for it to remain a credible emergency backup for UQADM.

The described process could conceivably have been quite arduous, but in this case it was not so difficult, since the options open to us were quite limited.

Current Load and Growth Estimates

User Numbers

In November-December 1992, it was normal for there to be in excess of 125 interactive users logged into UQADM during the day. Due to the commissioning of a new finance system in February, 1992, it is expected that there will be a growth of approximately 30 interactive users early in the year, although this increase will be reversed late in the year when the old finance system is decommissioned. It is also expected that there will be slow growth in usage of the student system by course advisers, faculty and departmental administration officers, and AIS staff. It is reasonable to expect that the interactive process load, in terms of both memory and processor resources, will grow by approximately 25-30% in 1993.

Mass Storage (Disk)

Administrative computing, as at December 1 1992 had sole use of 9 disk spindles with a total capacity of 12.2 Gigabytes of storage. Allowing 500 MB for paging and swapping space, a further 200 MB per spindle for housekeeping, and 1GB for scratch space leaves approximately 9GB for applications and data files. The total size of primary database files as at December 1 was 7.2GB, leaving just 1.8 GB for executable images and other necessary files. In other words, at the end of 1992, administrative applications were using all available mass storage. (In fact, some storage had to be rented from Prentice to cope with overflow.)

There will be considerable growth in storage needs in 1993. In addition, the total space for paging and swapping as at December 1 was inadequate, as was the distribution of files across disk spindles. Estimates of **growth** in storage needs for 1993 are:

System	Estimated Growth (MBytes)
Workorders	65
Equipment Register	13
Parlairs	400
Financial System (old)	500
Financial System (new)	2000
New Convocation Roll	260
Data Dictionary, Index files	25
Student System	750
Staff System	900
DB Admin Work Area	1250
Pagefile/Swapfile	500
Catch-up 1992 requirements	3500
Return space rented from Prentice	250
Total expansion 1993	10420

Provision of this space would, by virtue of its adding new spindles, also allow for better distribution of files and paging space to avoid some performance bottlenecks.

Mass Storage (Tape)

The increase in storage will put further pressure on backup systems. It is necessary that the entire disk farm be able to be backed up in four nights. Otherwise, backups will need to take place in prime time, with associated severe performance degradation.

Printing

The decommissioning of the DEC-10 system and its printers means that all Administrative printing is now performed on the lineprinters attached to the VAXCluster. Since the current printing capacity was not planned for this load, there are now constant printing backlogs. In addition, the monthly duty cycle of the printers is being exceeded by a factor of almost 2, causing frequent breakdowns and hence even longer queues.

Current lineprinting load is approximately 160 000 sheets per month. Current capacity is approximately 100 000 sheets per month.

Standby Processor

At present, the administration computer (UQADM) is backed up by a compatible, but smaller, processor (UQVAX). Currently UQVAX is maintained at a level of half the processor capacity and half the memory capacity of UQADM in order for it to represent a credible backup capability. If this policy is maintained, then UQVAX will need to be upgraded in line with upgrades to UQADM.

Development System

Another issue involves current undesirable practices forced by mass storage shortages.

- Administrative Information services performs application development in a separate environment provided by a dedicated VAX/VMS system. However, due to storage constraints, it is not possible to provide separate development database instances in line with normal industry practice. Much testing is forced to take place using production databases, with the concomitant risk of corruption and performance problems. It would be desirable to make a spindle available to the development system solely for development testing purposes.
- The development system is currently a member of the VAXCluster, along with UQVAX and UQADM. This causes extra load on the cluster system disk, which is currently the most heavily-used spindle.
- The fact that the development system is a member of the cluster and in addition uses a number of production database files means that it is currently infeasible to test new releases of system software for compatibility with administrative applications.

Results and Recommendations

Memory

As at the end of 1992, UQADM had 256 MB of memory installed. Even under existing load, there was very significant paging and swapping activity resulting in high demands on disk devices and poor to very poor response to user queries. Analysis of memory usage using simulation shows that another 128 MB of memory is required just to satisfy current requirements, and that a further 128 MB will be required in the course of 1993 to cope with the expected 25-30% growth in the interactive user load.

Recommendation: purchase 256 MB of memory for UQADM.

COST: \$148 000

Filestore (Disk)

As can be seen from the table above, approximately 4GB of capacity are required to upgrade storage immediately, 2GB is required for the new finance system, and another

4GB to allow for expansion in 1993. This makes a total of 10 GB, or 5 MDI230 winchester drives (or equivalent.).

Recommendation: purchase 5 Micro Tech MDI 230 winchester drives or equivalent.

COST: \$58 000

Examination of performance data gives some cause for concern that some disk channels may become saturated. One way to avoid this would be to upgrade the channel controller from an HSC50, which does not support caching, to an HSC95 which does. In particular, this would improve system disk performance.

Recommendation: upgrade HSC001 from an HSC50 to an HSC95.

COST: \$39 500

Backup Devices (Tape)

There are two upgrade possibilities for the backup system. One is to install a 42 GB cartridge loader system. A less expensive option would be to upgrade the current 2.3 GB 8mm drive to a 5 GB unit.

Recommendation: Upgrade current 2.3GB unit to a 5GB unit.

COST: \$16 000

OR:

Purchase a DEC TA867-AB 42GB cartridge tape loader.

COST: \$32 500

Processor

As at '0' Week 1993, the extra 256 MB of memory has been installed and CPU utilisation due to interactive users rarely rises above 60%. Since Administrative Information Systems staff are currently having some success in tuning applications, and since total cycle usage including batch processes over 24 hours is well under 50%, it is almost certain that a CPU upgrade will not prove urgent in 1993. Current developments in computing architecture also suggest it is prudent to wait until 1994 before spending money on processor capacity.

Should there prove to be performance problems despite the following of the other recommendations, there are a number of cheaper measures that could be taken which, together, would improve performance at least as much as the addition of another CPU.

- Installation of hardware shadowing on the system disk would solve incipient I/O saturation problems. This option would require the purchase of a disk and the necessary controller and software, at a cost of approximately \$15-20 000.
- The version of RDB currently in use is V3.1a. This version is so old that it is now unsupported. It should be upgraded as a matter of urgency.
- Redevelopment of AIS applications to reduce the number of attaches/detaches to databases and similar improvements have the potential dramatically to reduce contention for certain key files, as well as encouraging users to abandon the current practice of maintaining multiple sessions to use different application options.
- If necessary, computing time can be purchased on the UQVAX system. Batch processes, for example, could be submitted to UQVAX transparently to the user.
- All indications point to I/O bandwidth being a problem some time before CPU cycles are exhausted. The purchase of, for example, a HSC95 I/O processor for the 'hottest' disk units would alleviate this problem.

Recommendation: Tune applications, make use of batch facilities.

COST: \$0

Printer Capacity

Immediate needs, both in terms of raw capacity and of duty cycle, can be met by the purchase of two second-hand LP29 line printers. These printers would replace the current B1000 printers, which could either be sold, or, perhaps more advantageously, relocated to the J.D. Story building.

Recommendation: replace B1000 printers with second-hand LP29 printers.

COST: \$16 000

Recommendation: relocate B1000 printers to JD Story building.

COST: \$4500

Another option would be to reduce the impact printing load by making more use of laser printing facilities. This would also have the benefit of reducing the need for special stationary. It must be noted, however, that whereas impact printing costs about 1.5 cents per sheet, laser printing costs about 3.5 cents per sheet if the printer is used at its maximum duty cycle of 70 000 sheets per month. In the following estimate, it is assumed that about 70 000 pages of printing per month would be moved from impact to laser printing. This would represent the cheapest possible option, since this is the maximum duty cycle of the AIS LPS40. This reduces the impact printing load below the recommended duty cycle of the existing printers, and hence no printers would need to be purchased.

Recommendation: Move 70 000 sheets of printing per month to the existing LPS40+ laser printer, obviating the need to purchase new printing capacity.

COST: approx. \$10 000 p.a.

Backup Processor

If the current situation where the UQVAX processor is maintained at approximately half the capacity of UQADM is followed, the UQVAX will need a memory upgrade as well.

Recommendation: purchase 128MB of memory for UQVAX.

COST: \$74 000

Development System

It is a matter of some urgency that the development system be supplied with enough disk capacity to carry test versions of administrative databases.

Recommendation: purchase a 1.3GB disk drive for UQEVE

COST: \$5 000

Further, it is highly desirable that UQEVE be separated from the production VAXCluster so that new releases of system software can be tested and to reduce the I/O load on the VAXCluster system disk. This will entail the purchase of a further 1.3GB disk for UQEVE.

There is a possibility that it will also require the purchase of a tape drive for UQEVE, although SLS documentation suggests that UQEVE would still be able to be backed up across the network by the cluster.

Recommendation: Remove UQEVE from the VAXCluster.

COST \$5-10 000

Ian Burgess

Mark Williams

APPENDIX: Performance Information

The following pages contain a number of graphs showing performance indicators for the 10th of March, which was a typical Wednesday for UQADM Processing. These graphs are included, partly to show the grounds for the conclusions of this report, and partly as information for AIS staff.

CPU Statistics

The first two graphs show CPU utilisation for UQADM and UQVAX. As can be seen, there is a good deal of room for growth. At absolute peak times, it would require a growth of about 30% in interactive load before full capacity were reached. The batch load should not affect users' access to CPU, although it does cause some disk contention if disks are heavily loaded. Total 24-hour CPU Utilisation on UQADM was approx. 50%, with an 80% load between 0800 and 1800.

UQVAX also has a good deal of available CPU resources, so the daytime batch load could be shifted to UQVAX if CPU utilisation were to become a problem.

At current growth rates, there should be no reason to add CPU capacity within the next 12-18 months.

The next 2 graphs shows statistics for the number of processes. It is interesting to note that there is a peak process count of nearly 500 processes, of which about 400 were user processes. This process count is so high partly because many users are typically logged in several times to access different options in the same application. If this practice could be made unnecessary, the number of users could grow without any growth in the process count or demands on memory. At no stage were any processes swapped out. This supports the conclusion that the memory problems were largely solved by the purchase of 256 MBytes of memory.

Memory and Paging

As can be seen from the next 2 graphs, the amount of memory is now more than sufficient for current requirements. The near total absence of hard paging and the absence of swapping are proof of this.

Disk and File Statistics

The following 5 graphs will be of use to database administrators in dividing up the disks that have been ordered and in the partitioning of database files.

HSC Statistics

The remaining graphs show that neither of the HSCs are even approaching their capacity with respect to either throughput or to I/O operation numbers.

Other analysis does, however, show, that there is some cause for concern with individual channels and requestors. If these problems are not alleviated by load distribution when the final 4 disk spindles arrive, then there may be some justification for the upgrade of HSC001 from an HSC50 to an HSC95, since this allows write-through caching and reduces the load on channels and disks.

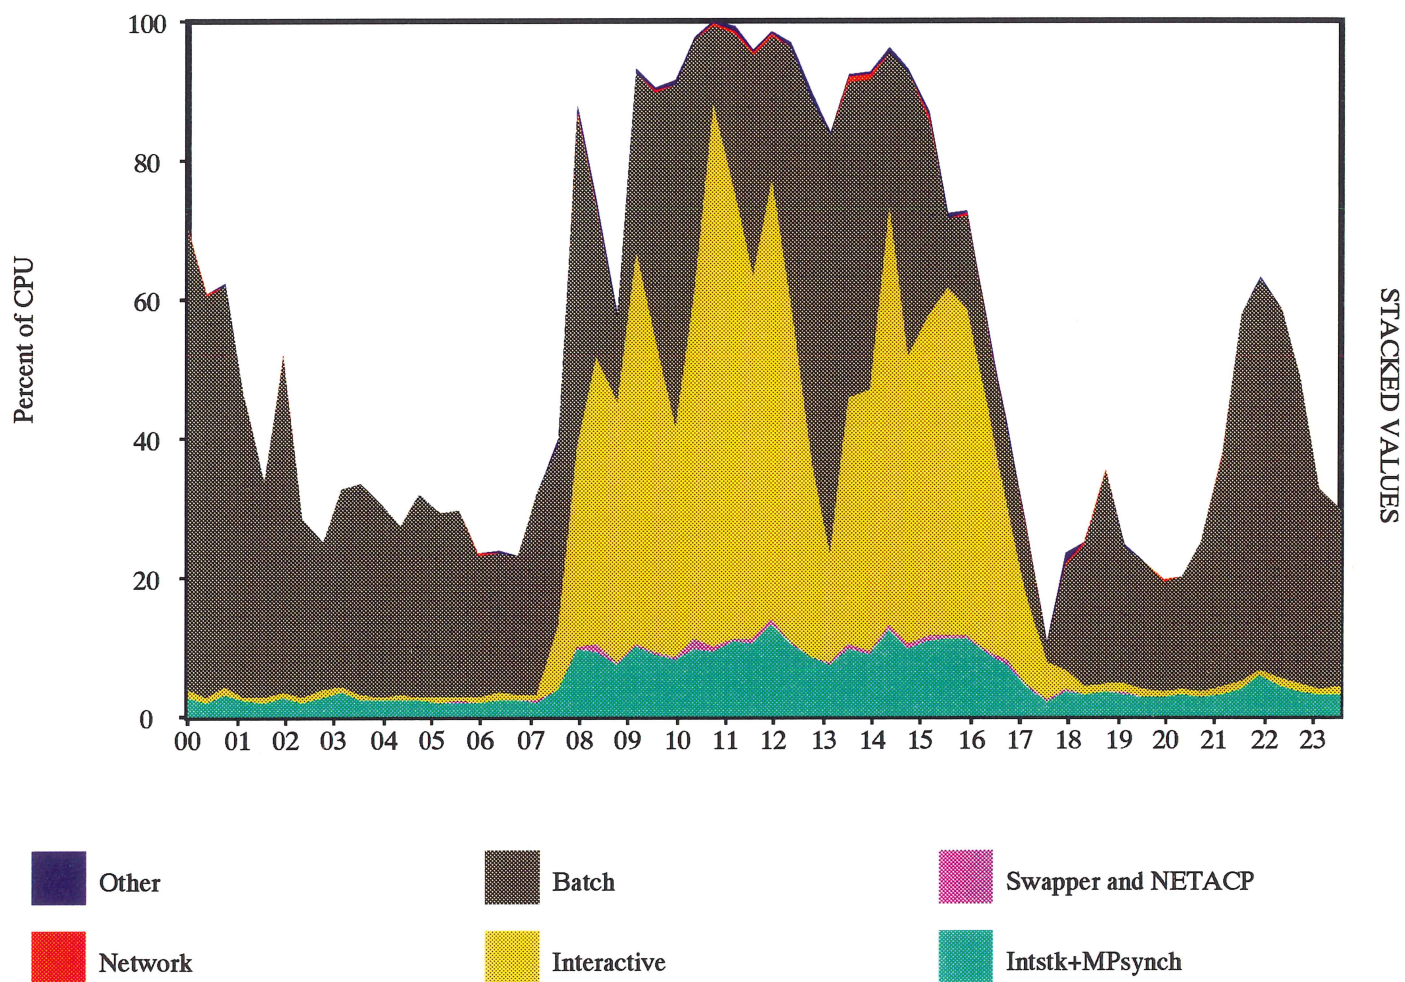
DECps

V1.0

CPU UTILIZATION

Node: UQADM

Date: 10-MAR-1993 00:00-23:58 (60 Pts Plotted)



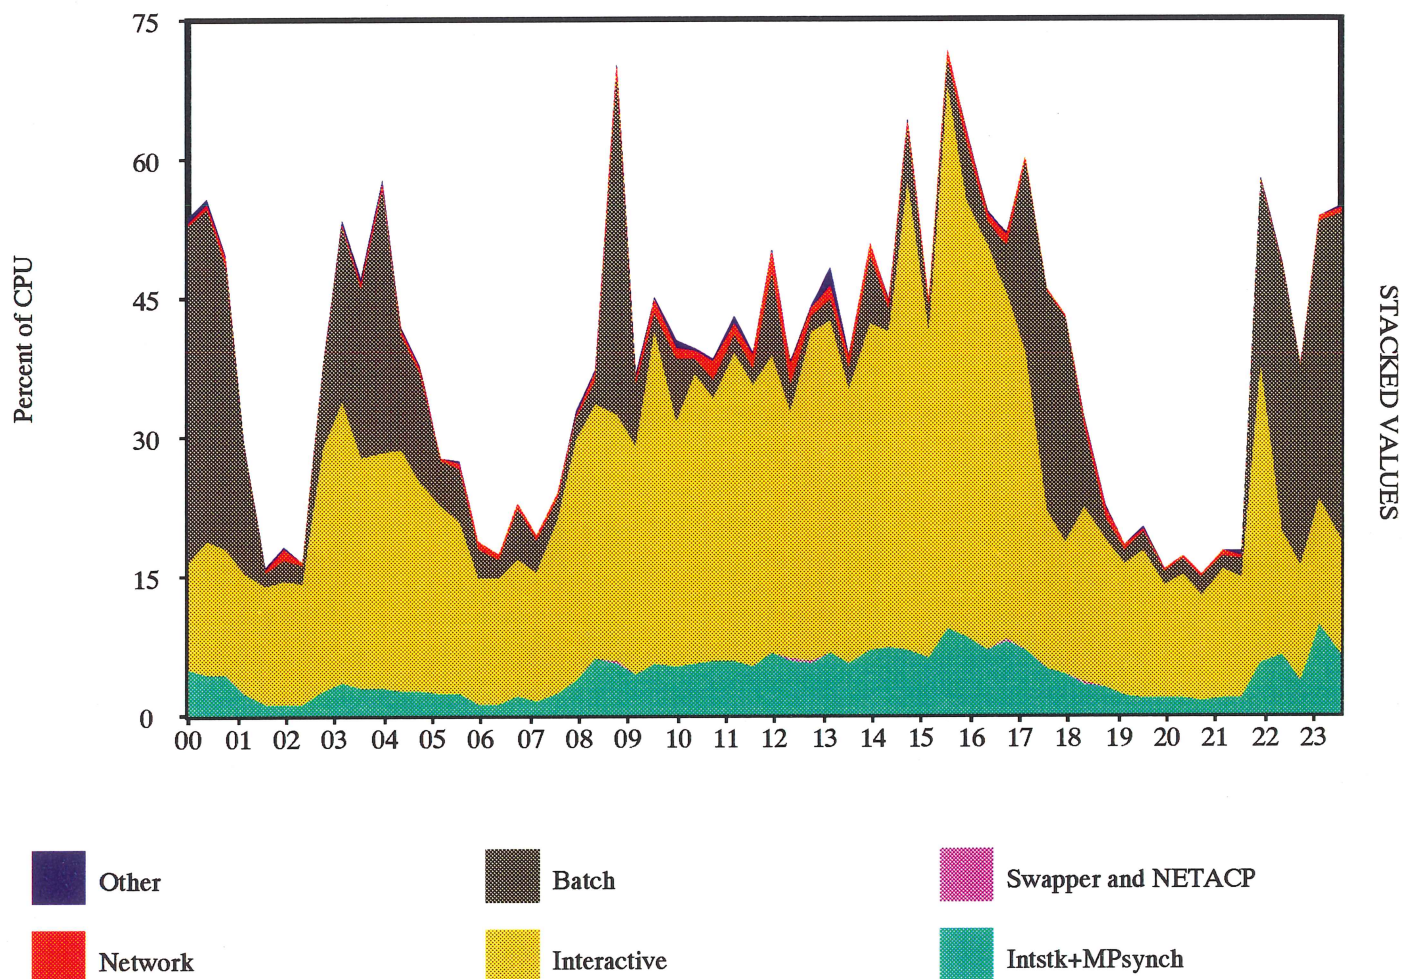
DECps

V1.0

CPU UTILIZATION

Node: UQVAX

Date: 10-MAR-1993 00:00-23:58 (60 Pts Plotted)



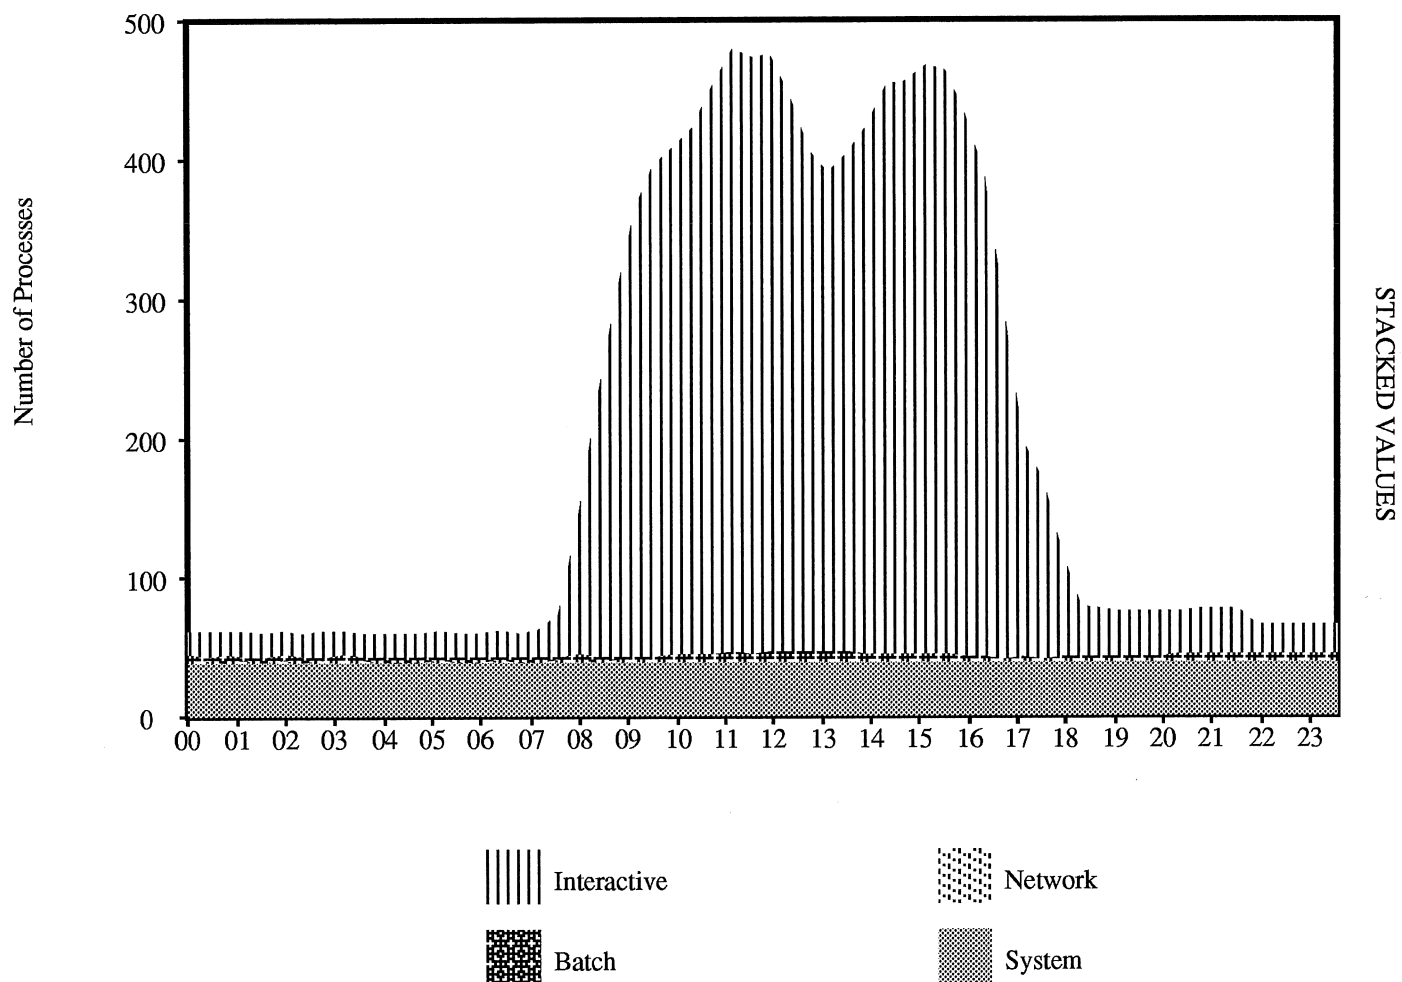
DECps

NUMBER of PROCESSES by Type

V1.0

Node: UQADM

Date: 10-MAR-1993 00:00-23:58 (60 Pts Plotted)



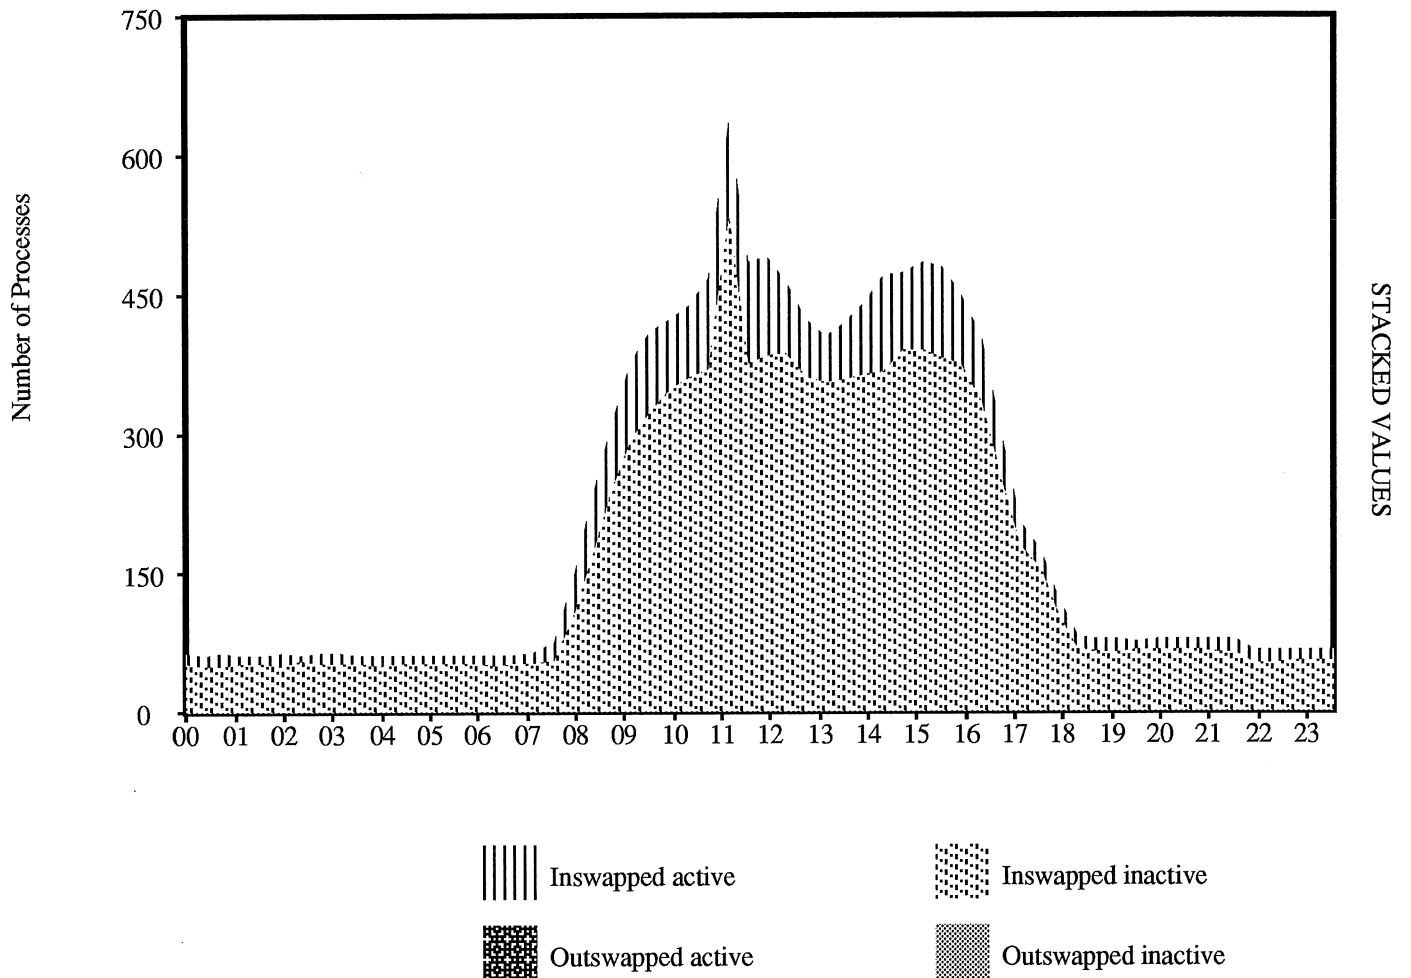
DECps

NUMBER of PROCESSES by State

V1.0

Node: UQADM

Date: 10-MAR-1993 00:00-23:58 (60 Pts Plotted)



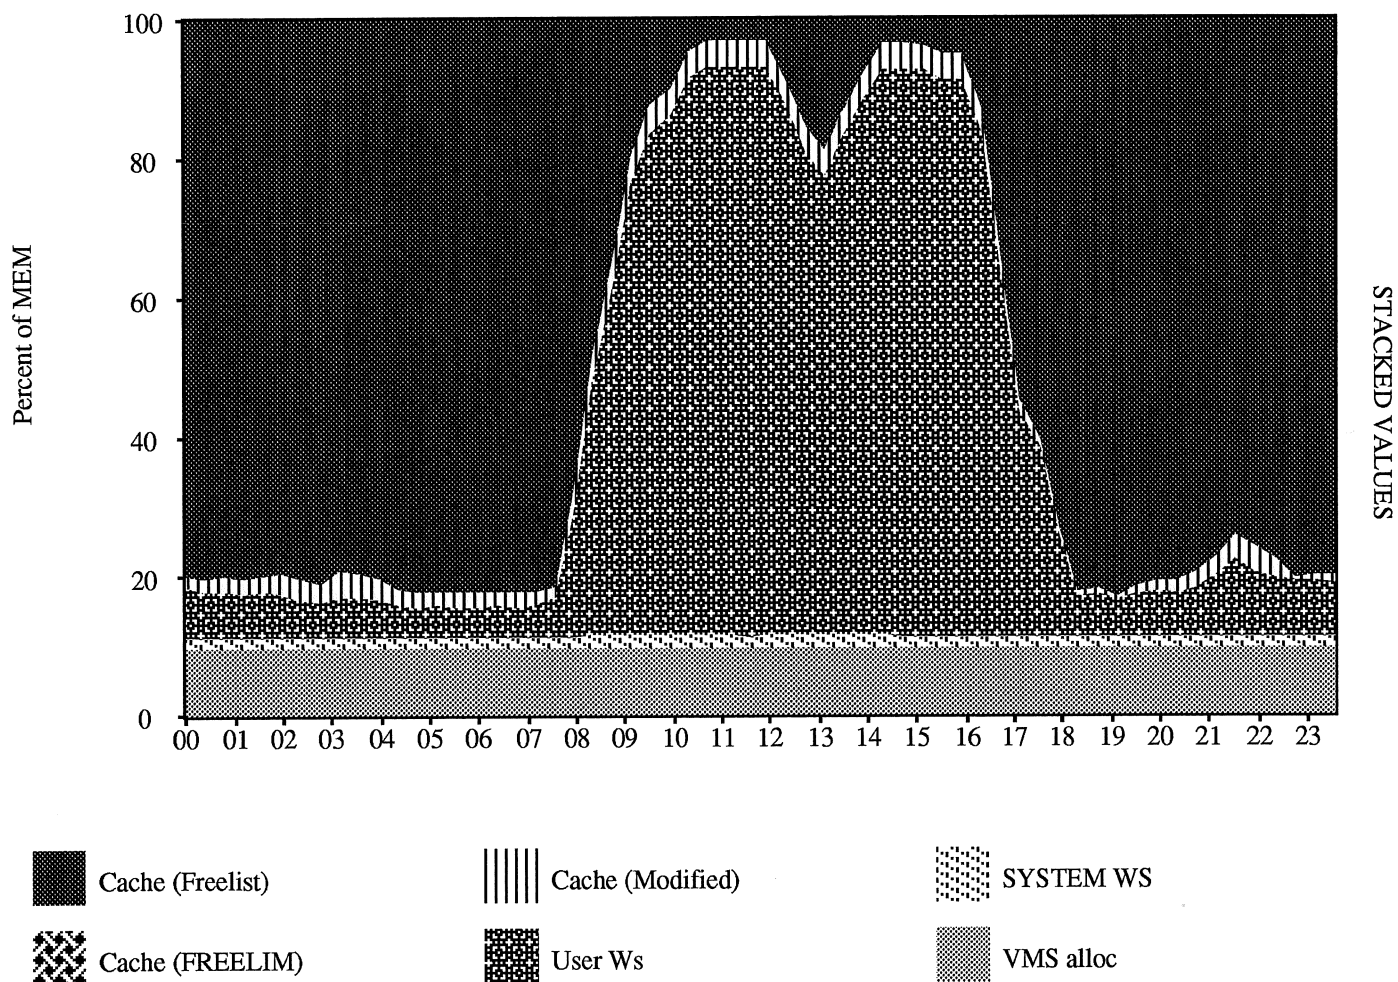
DECps

V1.0

MEMORY UTILIZATION

Node: UQADM

Date: 10-MAR-1993 00:00-23:58 (60 Pts Plotted)



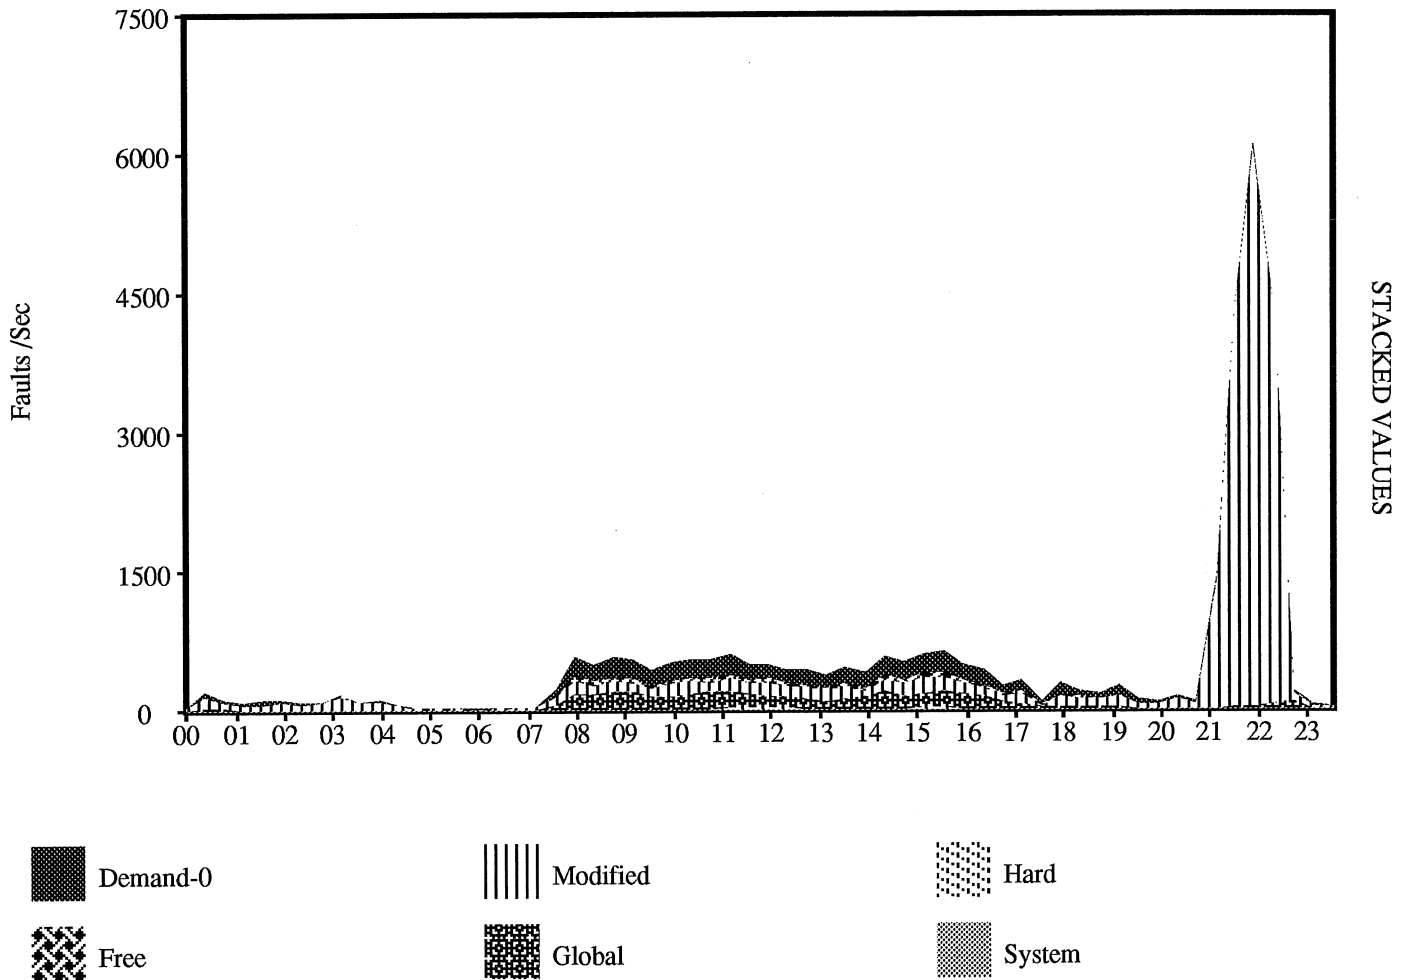
DECps

V1.0

PAGE FAULT RATE

Node: UQADM

Date: 10-MAR-1993 00:00-23:58 (60 Pts Plotted)



DECps

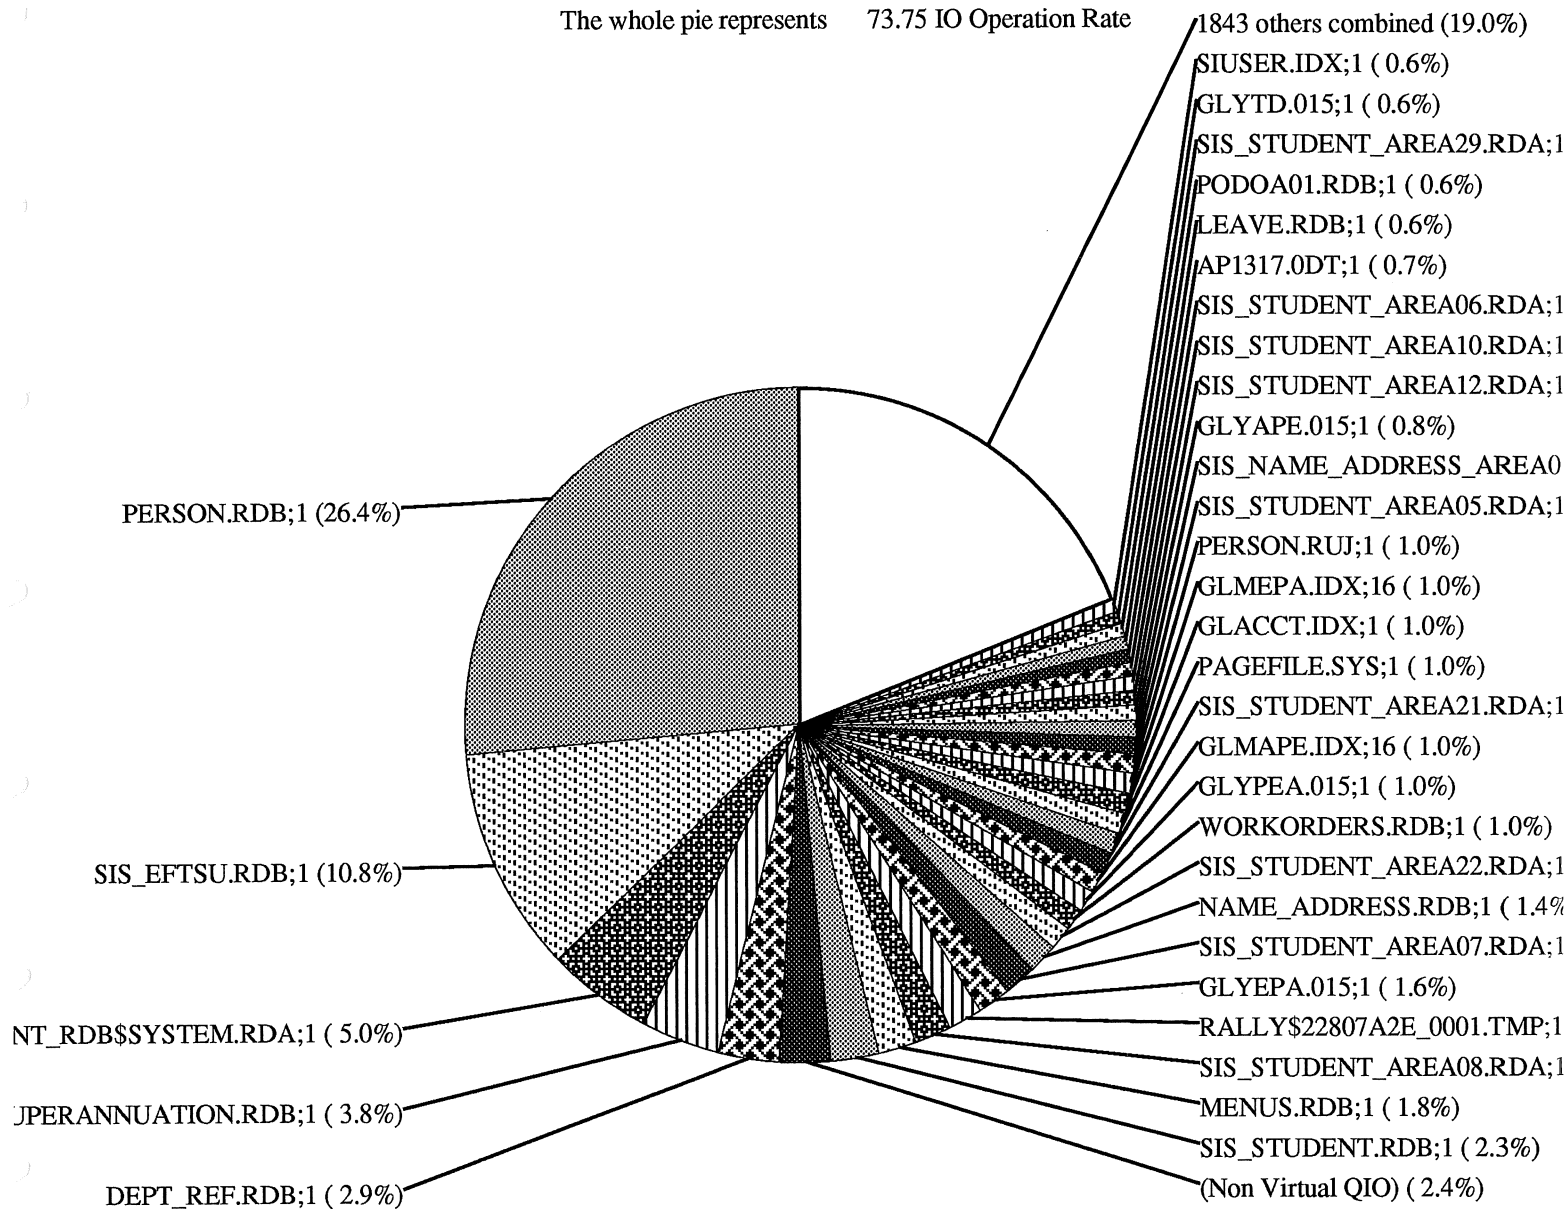
V1.0

TOP OPERATIONS FILE

Node: UQADM

Date: 10-MAR-1993 00:00-23:58

The whole pie represents 73.75 IO Operation Rate



DECps

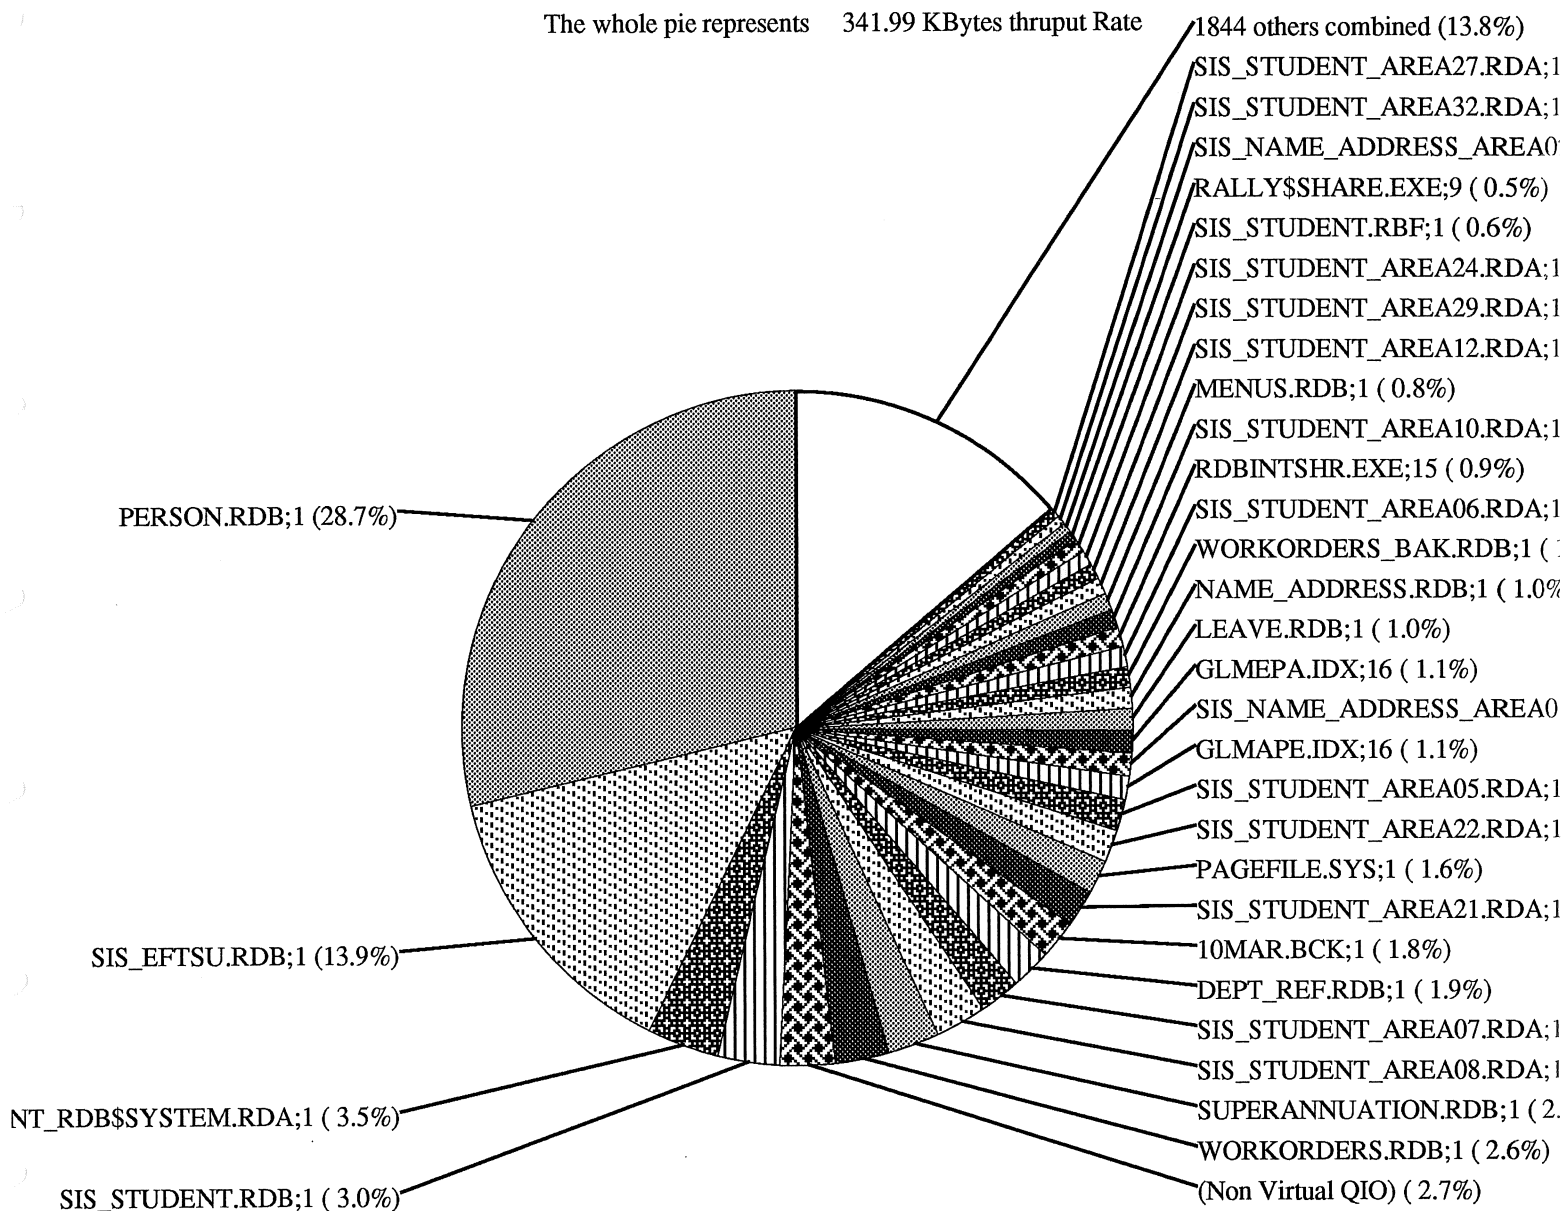
V1.0

TOP THROUGHPUT FILE

Node: UQADM

Date: 10-MAR-1993 00:00-23:58

The whole pie represents 341.99 KBytes thruput Rate



DECps

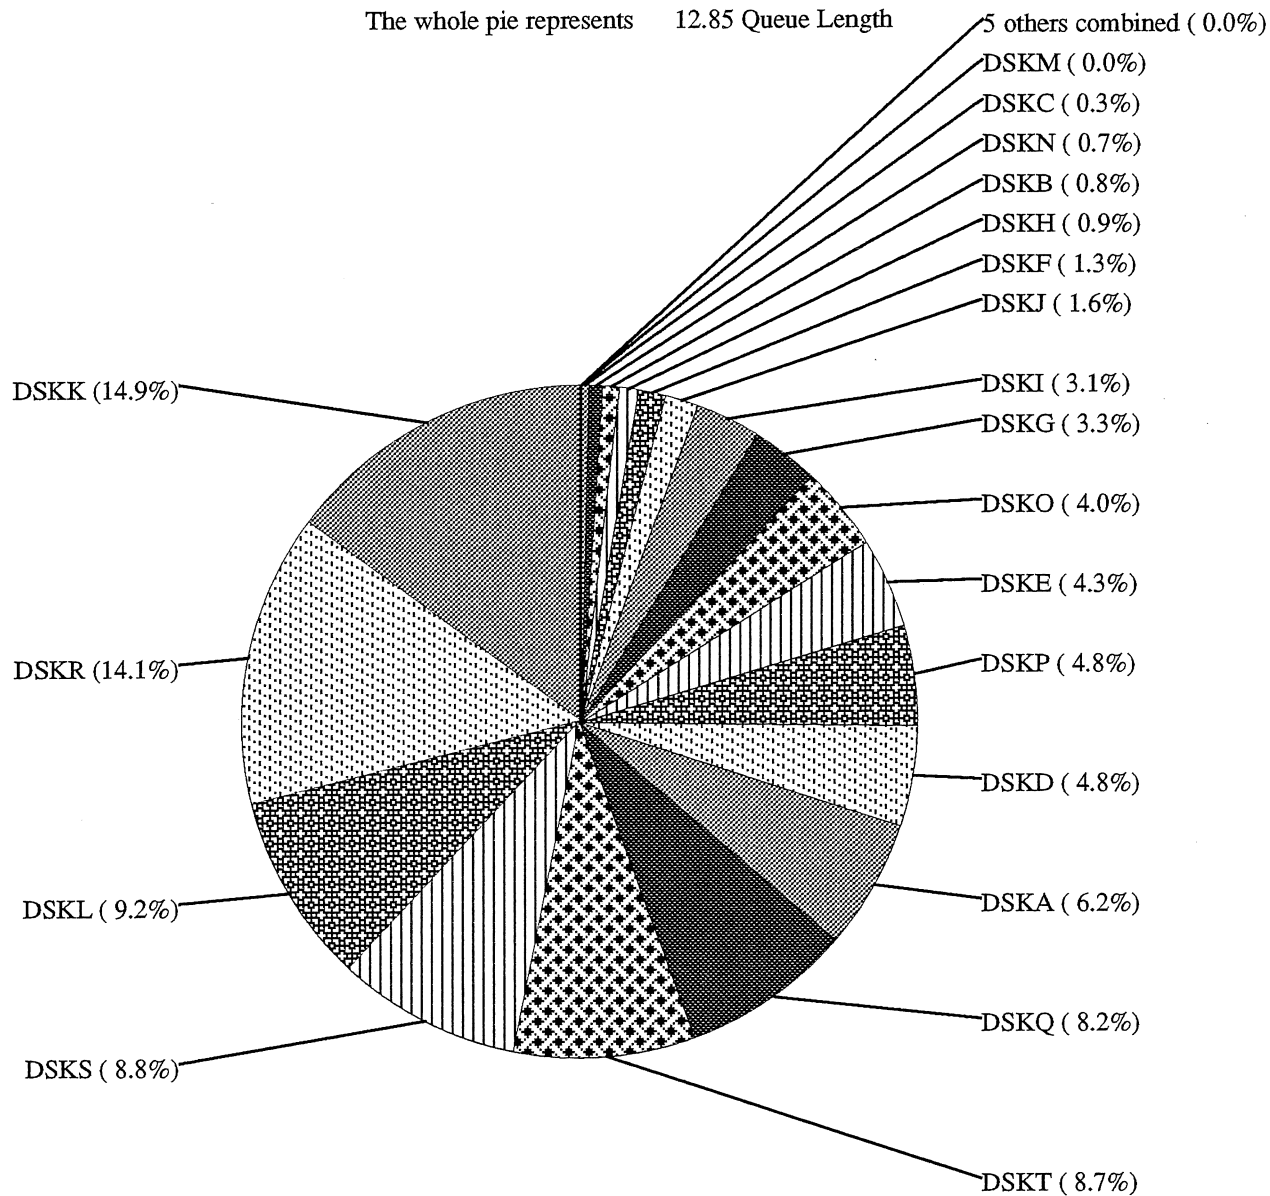
V1.0

TOP QUEUE DISK VOLUME

Nodes: UQVAX, UQADM, UQEVE, EAGLE

Date: 10-MAR-1993 00:00-23:58

The whole pie represents 12.85 Queue Length



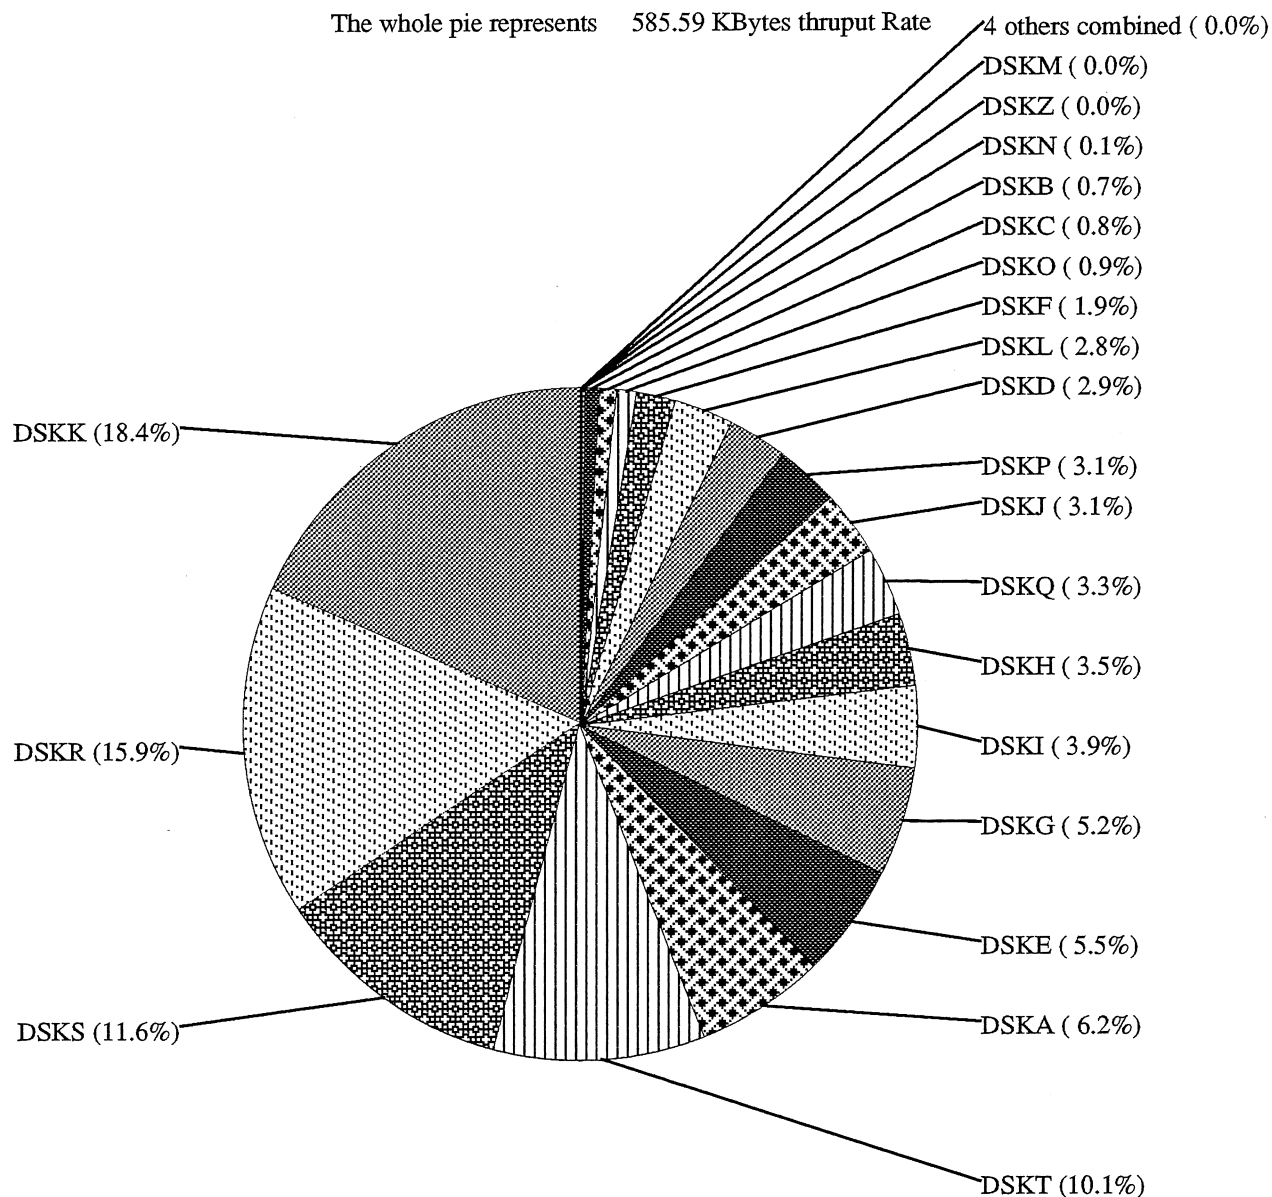
DECps

V1.0

TOP THROUGHPUT DISK VOLUME

Nodes: UQVAX, UQADM, UQEVE, EAGLE

Date: 10-MAR-1993 00:00-23:58



DECps

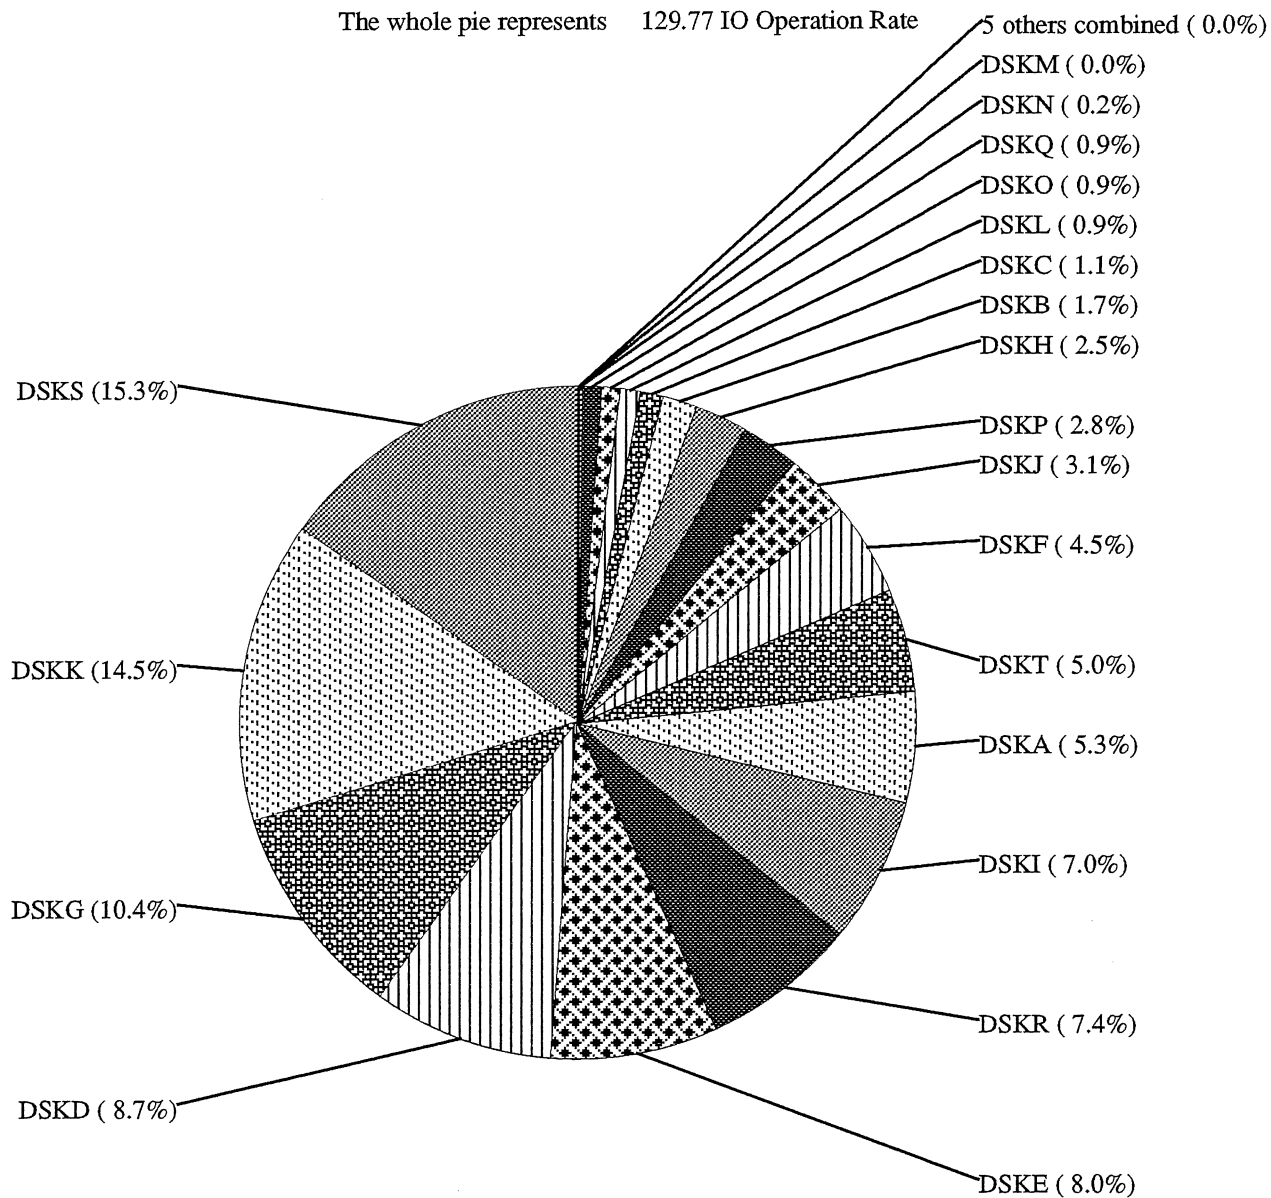
V1.0

TOP OPERATIONS DISK VOLUME

Nodes: UQVAX, UQADM, UQEVE, EAGLE

Date: 10-MAR-1993 00:00-23:58

The whole pie represents 129.77 IO Operation Rate



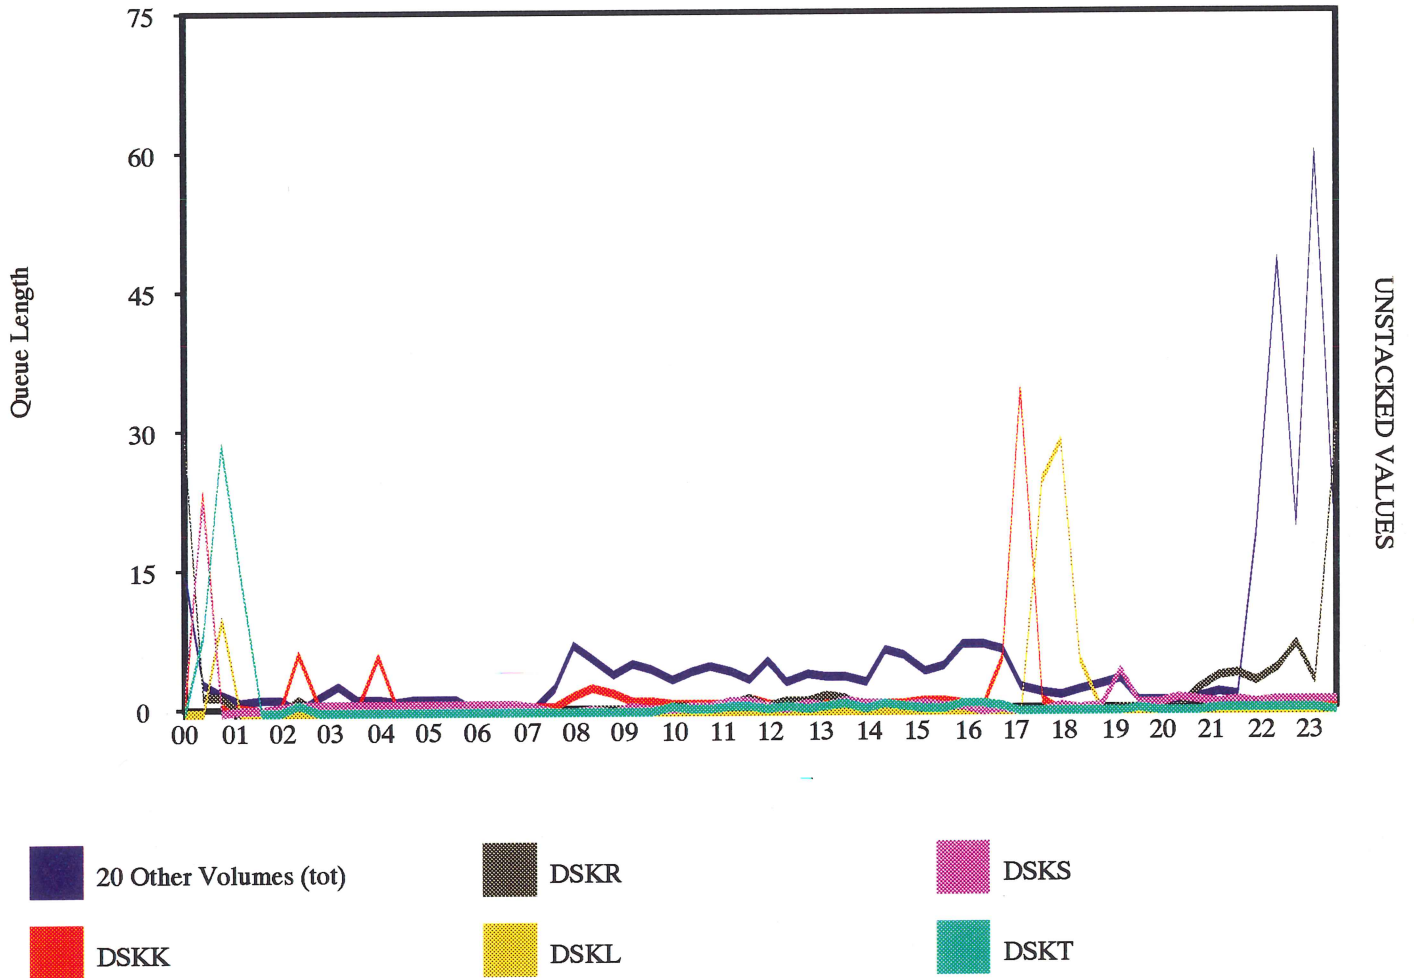
DECps

V1.0

TOP QUEUE DISK VOLUME

Nodes: UQVAX, UQADM, UQEVE, EAGLE

Date: 10-MAR-1993 00:00-23:58 (60 Pts Plotted)



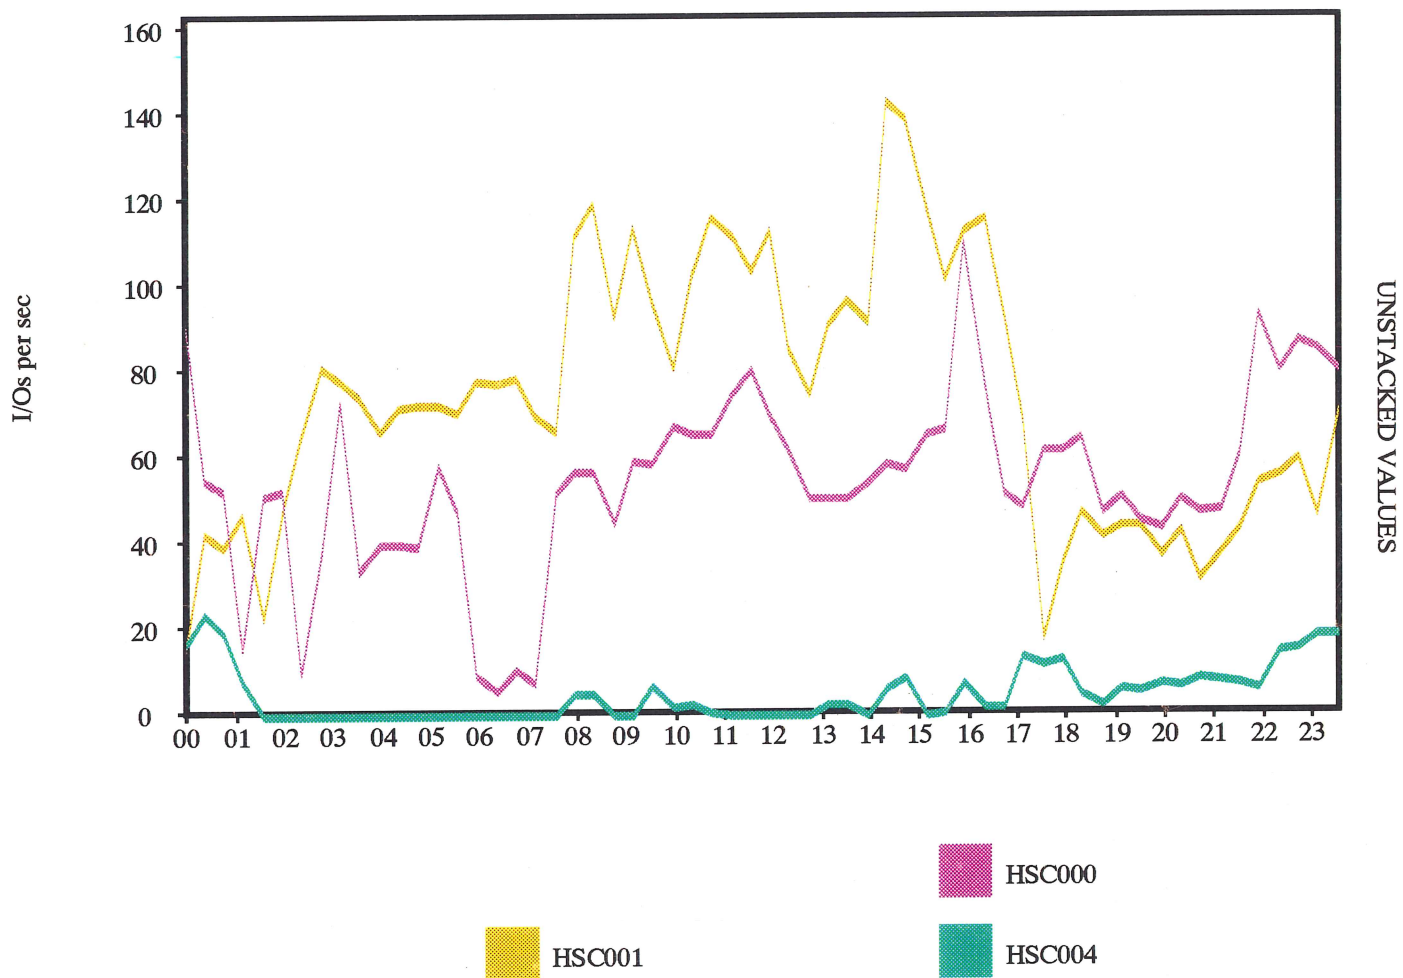
DECps

V1.0

TOP HSC IO

Nodes: UQVAX, UQADM, UQEVE, EAGLE

Date: 10-MAR-1993 00:00-23:58 (60 Pts Plotted)



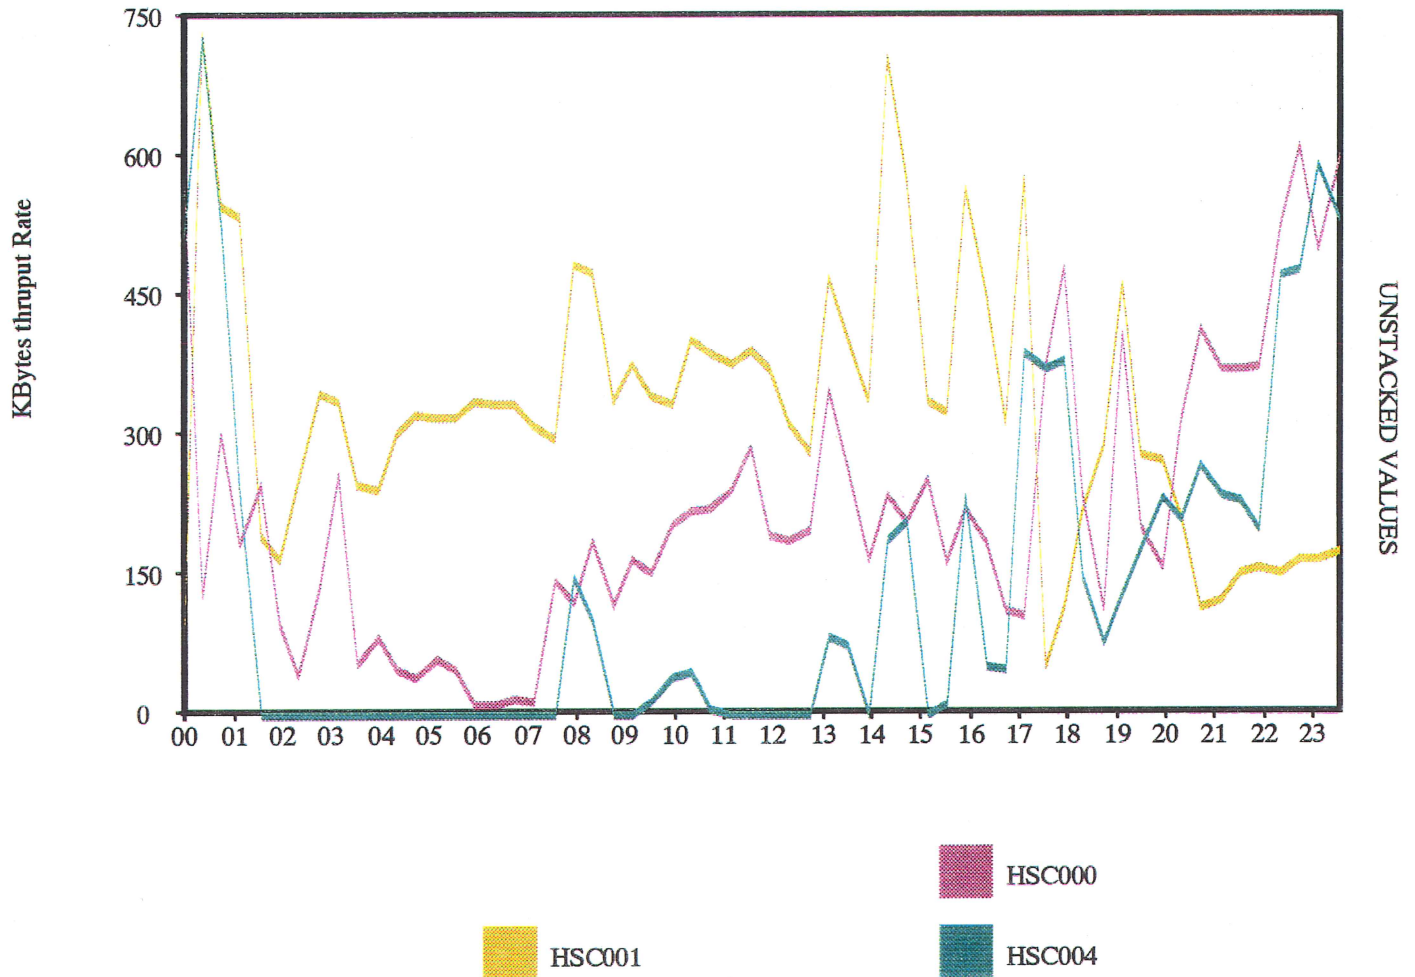
DECps

V1.0

TOP HSC THRUPUT

Nodes: UQVAX, UQADM, UQVEVE, EAGLE

Date: 10-MAR-1993 00:00-23:58 (60 Pts Plotted)



DECps

V1.0

TOP QUEUE HSC CHANNEL

Nodes: UQVAX, UQADM, UQVEVE, EAGLE

Date: 10-MAR-1993 00:00-23:58 (60 Pts Plotted)

