ICANS-XI International Collaboration on Advanced Neutron Sources
KEK, Tsukuba, October 22-26, 1990

Future Strategy for Computing at ISIS

W.C.A.PULFORD
ISIS Instrumentation Division,Rutherford Appleton Laboratory

ABSTRACT

This paper initially provides a brief description of the current system together with some of the operational problems encountered during the five years of ISIS running. It then summarizes some of the thoughts of the ISIS computer and electronics groups concerning the best ways for the ISIS data acquisition and analysis systems to evolve. Topics considered include: the increasing use of workstations and windows, the possible future use of UNIX and the next generation of data acquisition electronics.

I am much indebted to the whole IID computing group (Kevin Knowles,Chris Moreton-Smith,Tony Valente,Richard Brodie,Kate Crennell and Scott Robertson) as well as Julian Norris and John Holton of the electronics group, for their contributions to this paper.

Introduction

A powerful and successful pulsed neutron source such as ISIS must be supported by a computer system of high capability. It is clearly of limited value to run this powerful source without the means to progress rapidly and easily from raw data to scientifically interpretible results. The PUlsed Neutron Computer Hierarchy or PUNCH computer system is used for all collection and subsequent analysis of data collected from the neutron diffraction instruments arrayed about the ISIS pulsed neutron source. At present it consists of a Vax8650 computer system, the HUB, connected in a Local Area Vax Cluster (LAVC) to 9 x VaxStation 3200, 6 x MicroVax II, 1 x Vax750 and 3 Vax730 computers. The smaller computers are known as Front End Minicomputers (FEMS) and are responsible for the data acquisition from the ISIS instruments. In addition there are 2 VaxStation 3100, 2 VaxStation 2000 and 12 IBM PC workstations as well as about 130 dumb visual display terminals connected to the network of computers. The whole PUNCH system is displayed schematically in figure 1.

The Data Collection Process

ISIS is a pulsed source of both neutrons and muons. In the former case, single or multidetectors are used to count neutrons scattered by a sample of the material under investigation. In the latter case positrons produced by muon radioactive decay within the sample are counted using multidetector arrays. For both the pulsed nature of ISIS confers a time structure on the detected events. This time structure is reflected in the data collected by storing the counts as a function of their time of occurrence in time bins of width δt (typically 1 μs for neutrons and 20 ns for muons). Consequently a histogram of data is built up for each detector element. These histograms are stored in computer data files whose size is dictated by the product of the number of time bins used and the number of detector elements. Each ISIS instrument is optimised for the investigation of particular types of condensed matter.The sizes and the rates of collection of these 'raw' data files varies considerably. Table 1 gives typical values of these parameters.

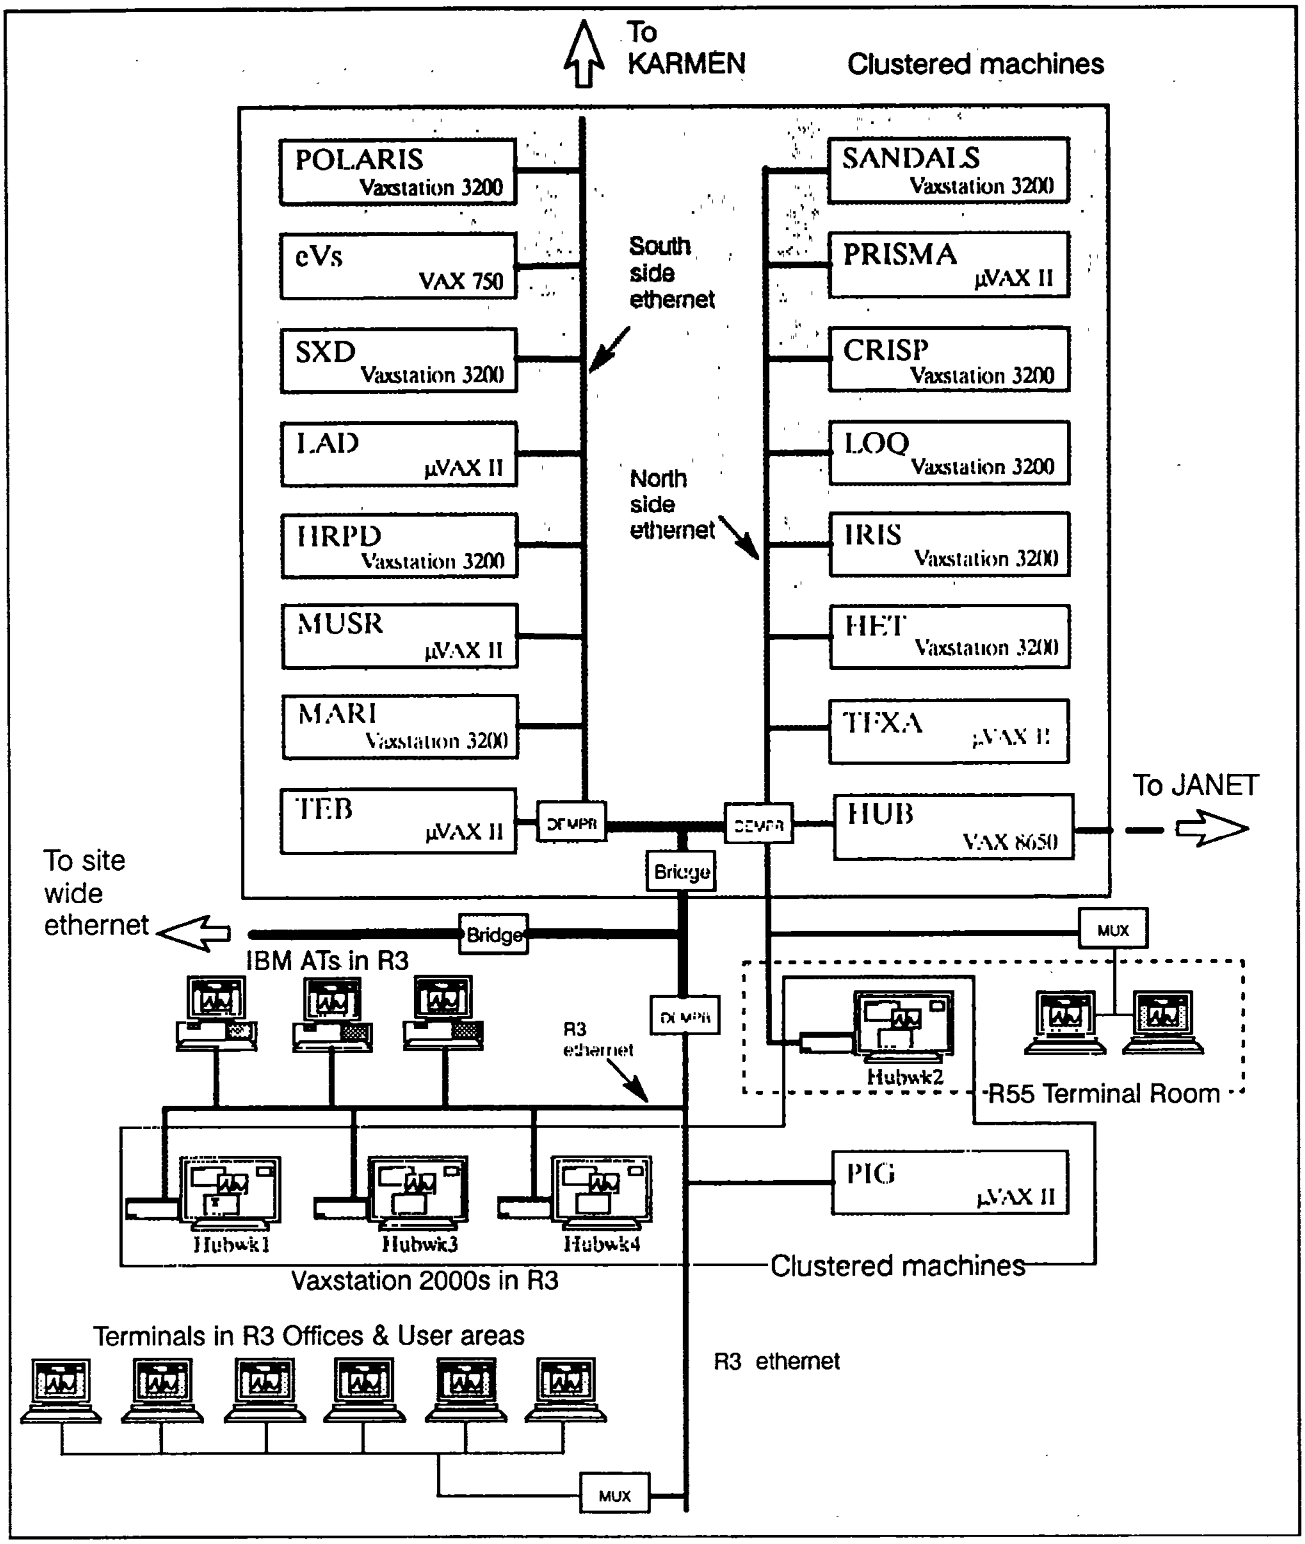

Figure 1 THE PUNCH COMPUTER SYSTEM

Instrument	Number of Runs	Average Data Volume (Kbytes)	Total Data Mbytes	Instrument type (E – Elastic, I – Inelastic)
CRISP	458	50	22	Surface Reflectometer(E)
LOQ	346	1200	415	Low Angle diffractometer(E)
POLARIS	245	330	80	Medium Resolution Powder Diffractometer(E)
HRPD	424	640	272	High Resolution Powder Diffractometer(E)
IRIS	137	180	25	Inelastic High Resolution Spectrometer(I)
LAD	99	790	79	Liquids and Amorphous Diffractometer(E)
PRISMA	153	186	29	Inelastic Excitations Spectrometer(I)
HET	70	1800	132	High Energy Transfer Inelastic Spectrometer(I)
SXD	149	1400	252	Single Crystal Diffractometer(E)
TFXA	30	3200	96	Time Focussed Crystal Analyser(I)

Table 1 – Data Collection Rates for Various Instruments for Cycle 89/7 (duration 4 weeks)

The purpose of the PUNCH computer system must be to archive these data files securely as well as facilitating their analysis to produce physically interpretible results.

Sequence Of Data Processing

Broadly, the sequence of data processing may be described as consisting of the following phases :

1.Setup

Optimum values for the experimental parameters such as the number of time bins and run durations are calculated and then set up for the data collection phase or "RUN". Also such physical parameters as the instrument geometry, sample orientation and temperature are defined at this stage.

2.Run

The progress of the data collection is monitored continuously during this phase by inspecting key run parameters and collection rates. This is important for the early detection of instrument errors or perhaps sample state changes.

3.End of Run

On completion of an experimental RUN the collected histograms are stored in the computer RAW data file. This file is automatically copied from local disk space onto an optical disk jukebox connected to the HUB computer. Details of the data archived in this way are stored in catalogue files for each instrument. This information is sufficient for the archived data to be restored to magnetic disk at a later date.

4.Data Reduction and Analysis

Data Reduction and Analysis – the processing of data varies considerably depending on the instrument and type of material under investigation although most uses a standard program environment called GENIE. A typical sequence of processing might be .

a) Merge histograms collected by detectors at equivalent positions in space such as those diametrically opposite at the same radius on a cylindrically symmetrical instrument.

b) Normalize data. This is basically a process of comparing the data against some collected in the absence of the sample to eliminate any systematic fluctuations in the instrument transmission.

c) Peak fitting. In most cases it is necessary to perform considerable analysis of the histograms in order to determine the exact positions and sizes of features. This may involve processor intensive least squares peak fitting routines, such as in RIETVELDT refinement, or as is becoming increasingly common, the technique of maximum entropy (MAXENT).

d) Model fitting. This depends totally on the type of experiment being performed and whether the probes used are neutrons or muons.

The first category of neutron instrument is that which concentrates on elastic scattering of the incident neutrons. On data collected from these instruments where the neutrons are scattered by the sample with no energy loss, the primary purpose is usually to establish the atomic or molecular structure of the material. In this case further processing involves proposing a model for the structure and using it to calculate the data points expected. These expected data are then compared with the observed data and a new refined model established by a process such as least squares (LSQ). The main line of this type of analysis is structure factor least squares which minimizes the discrepancy between the observed and calculated structure factors of the sample. (The structure factors (fobs) are closely related to observed intensities Iobs = |fobs|**2.) This method tends to be reliable but is a very heavy user of processing power scaling as the number of data observations and the square of the number of scattering elements in the sample. Monte-Carlo techniques tend to replace least squares in model fitting or generation for instruments investigating the properties of liquids and amorphous materials. This is a processor intensive method where the total energies are calculated for a large number of randomly generated models, the most likely model being indicated by that with the lowest energy.

The other major category of neutron instrument is that of inelastic where neutrons exchange some energy with the molecules or atoms of the material under investigation. The observed energy shifts then provide information about the atomic lattice motions or the motions of the atoms themselves. The analysis of this type of experiment tends to involve the calculation of many incident and resultant neutron energies and momenta. Also to appreciate the nature of the molecular motions it may be necessary to compare results obtained from a number of runs.

Muons are fundamental particles which resemble either heavy electrons or positrons depending on their charge. They are produced in large quantities at ISIS by inserting a thin target into the incident proton beam line. The general thrust of the experiments using these probes is to investigate diffusion effects within the solid state or to examine materials during magnetic state transitions, although more recently, considerable effort has been applied to examine the potential of muon catalysed fusion. In each case the computer processing involves the least squares fitting of various exponential functions to the observed data. The most accurate experiments are performed when each of the many detectors on a muon instrument is considered separately. Thus the processing involved scales linearly with the number of detectors.

5.Data Presentation

Once data analysis is complete it is necessary to present the experimental findings. Extensive use of desktop publishing and graphics is made to expedite this process.

The Functions of the PUNCH computers

The configurations of the current PUNCH computers are illustrated in Figures 2 and 3. The FEMs are used primarily for phases 1, 2 and 3 of the data collection process. The FEMs are connected in a peer to peer LAVC where each machine has a separate operating system. This, together with disk capacities ranging from 160 to 768 Mbytes, ensures data collection irrespective of the state of the HUB computer. Nevertheless the HUB computer does fulfil a pivotal role in the PUNCH computer system. This machine performs much of phases 4 and 5 in the data collection process, but its importance extends well beyond this. The HUB computer is the local point for data concentration and storage because it controls a large quantity of magnetic disk space and the 40 Gbyte optical disk jukebox. This latter storage ensures that any data ever collected from the ISIS instrument is available on line. The constantly expanding data storage requirements are illustrated by Figure 4.

It is evident that the data storage requirements continually increase. Indeed, in order to maintain the online status of ISIS data is has been necessary from Cycle 4 1989 to perform processor intensive file compression before archiving.

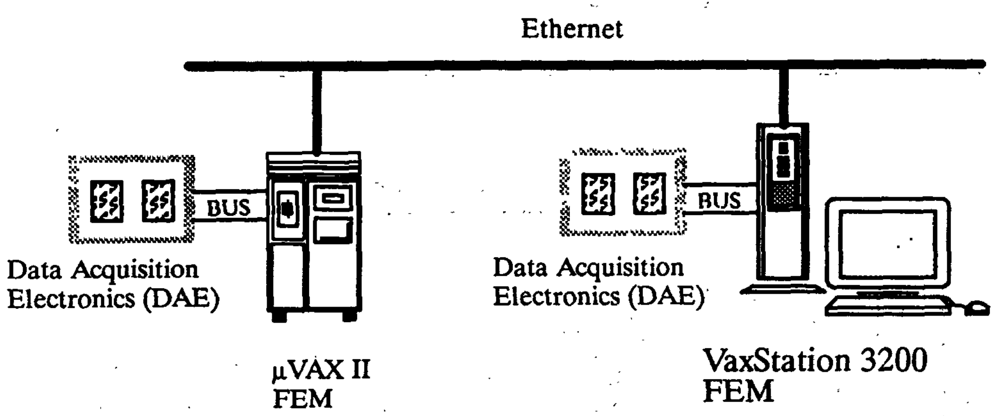

Figure 2 The Two Main Types of FEM

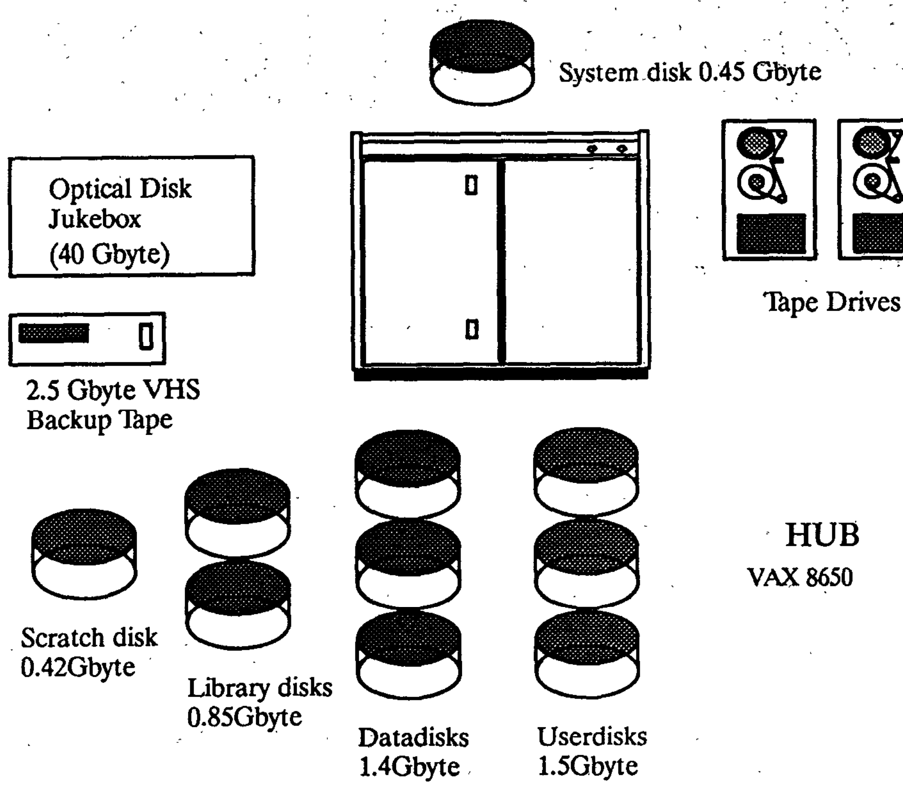

Figure 3 The HUB configuration

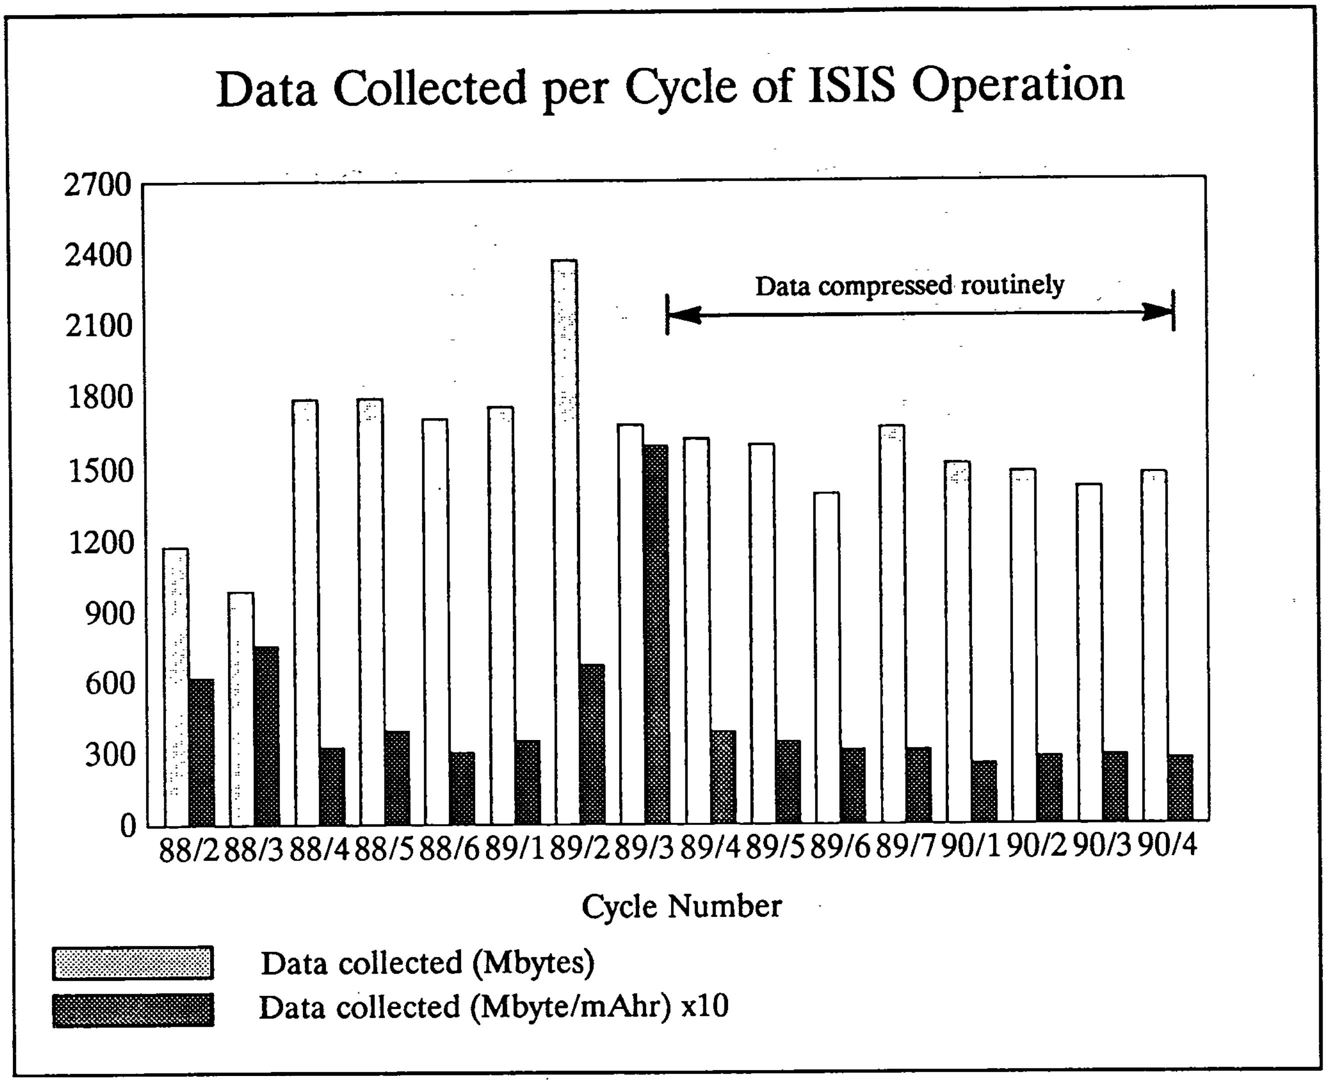

Figure 4

The volume of data collected per proton milliamp hour has remained approximately constant since Cycle 4 in 1989. This can be explained by the following observations:

a) ISIS has been running with a Tantalum target thus reducing the neutron flux by 50%.

b) Fewer more complex experiments are being performed which require longer to setup and longer run durations to improve statistics.

c)There have been no new scheduled instruments during this period. MARI and SANDALS are only now in cycle 90/5 making a significant contribution to the data volumes collected.

The HUB computer provides a broad spectrum of general computing for over 300 registered users. The users perform diverse sets of tasks ranging from editing and word processing through to data analysis often using graphical visualisation techniques. Figure 5 illustrates the typical workload of the HUB processor, the simulations (SI) are performed under the control of a low priority batch queue to use free processing power until pre-empted by any other process. Some reasons why the HUB workload has developed in this way are given below :

a) The HUB computer may perform processing without affecting data collection.

b) It is the only machine within the PUNCH system directly accessible from the UK academic network (JANET).

c) Until recently only the HUB had routine access to the ISIS data catalogues.

d) It is the most powerful computer in the PUNCH system.

e) Many software products are only available on the HUB computer because of their cost and/or storage requirements.

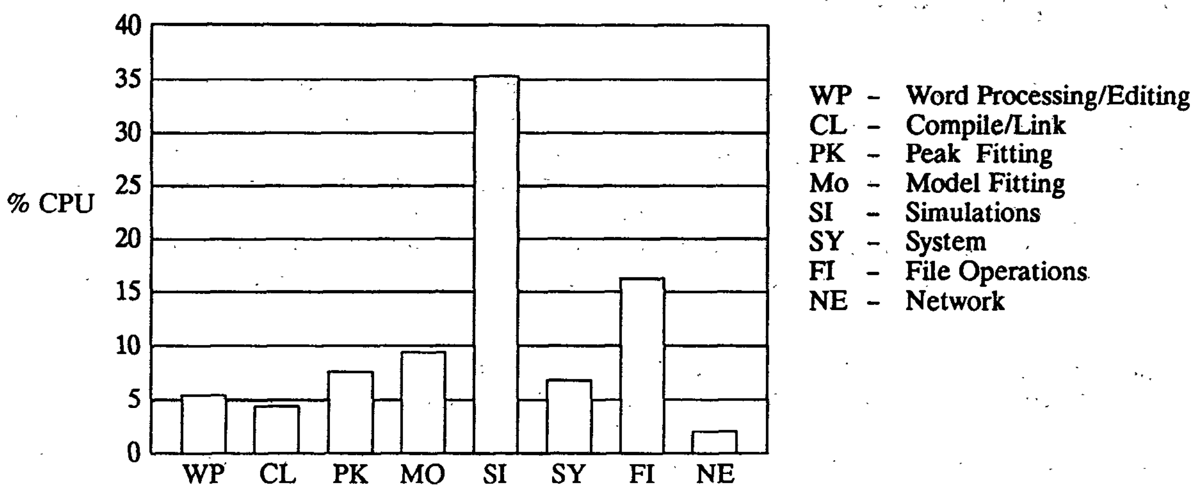

Figure 5 The distribution of HUB processor time between different categories of task

Current Problems With The System

Most of these are concerned with the HUB 8650 computer configuration. It has always been intended that this machine satisfy the interactive needs of a large number of users performing heterogeneous operations. The essence of an interactive machine must be responsiveness, particularly for editing and graphics. Subjectively the response to a typed command should take no longer than 2 seconds but this has deteriorated considerably at peak times over the past 6 months, despite the efforts described below to slow the effect. In order to maintain the interactive response, users are strongly encouraged to run processor intensive tasks in batch mode where they run at lower priority but which the operating system may schedule more efficiently. The scheduling has the unfortunate side effect that few or no batch jobs run during the peak hours between 0900 and 1700 hours daily. Another serious problem is that the HUB is chronically short of both optical and magnetic disk space. Currently the optical disk jukebox has about 8 Gbytes of its 40 Gbyte capacity free and with data accumulating at a rate of about 1.5 Gbyte per 6 week cycle we will no longer be able to maintain all data online from Cycle 6 of 1990. Magnetic disk space is also at a premium, particularly since the poor batch job turn-round tends to extend the duration of data files. By applying stringent disk quotas, the HUB runs constantly with its disks between 85 and 95% full with no spare capacity. In order to run more detailed data analyses the HUB users usually have to compete for the relatively small available space on a scratch disk. A further problem has been HUB downtime caused by a variety of hardware faults. The frequency of these faults can be ascribed to the fact that the 8650 uses 1984 technology which tends to use high current devices to obtain what is by current standards modest performance. This is reflected in the relatively high maintenance cost of about £30,000 per annum.

Remedial Actions

The past year has seen the progressive replacement of the original VAX 11/730 FEMs by VaxStation 3200 workstations. These have about 10 times the central processing speed and 5 times the disk capacity of the older machines. Unfortunately, these workstations, which actually run at about 45% of the 8650 processing speed, are restricted to a maximum of 2 interactive users and can therefore do little to alleviate the interactive load on the

PUNCH system. Nevertheless, connecting these computers in a local area vax-cluster does enable some of this processing power to be exploited by means of cluster-wide accessible batch queues. Any jobs placed in these queues are automatically run on any free cpu in the cluster. This has gone some way to increase the batch processing capacity of the PUNCH system, but it is subject to a law of diminishing returns. Each additional clustered computer does abstract a significant proportion of the HUB processing and disk capacity.

Future Considerations

It is our experience that the computing requirements generated by each instrument double every three years. This observation is shared with the our complementary institute, the Institut Laue Langevin at Grenoble. I believe this assessment of the situation at ISIS to be optimistic. As an example the facility has just received a grant of £2.4M to enable the doubling of the ISIS proton beam current to 200μA. This together with an imminent enhancement to the ISIS target will result in an up to fourfold improvement in neutron flux. Thus purely on the basis of intensity the number of experiments capable of being performed and hence the data collected may well more than double over the next 24-30 months. The problem of this impending data processing bottleneck will be exacerbated by two further factors:

a) Four new instruments MARI,ROTAX,SANDALS and eventually a second muon spectrometer will be commissioned. (MARI and SANDALS in 1990 , ROTAX in 1991 and the new MUSR in 1992)

b) Many of our current instruments (HRPD,POLARIS,CRISP,IRIS,LOQ and SXD) will be equipped with new large detectors with the potential to produce manyfold increases in the data volumes collected.

It is reasonable to assume that not all of this data volume and associated processing will be realized immediately. Much of the increased ISIS intensity will be used to provide better statistics or to enable the prosecution of more ambitious experimental projects possibly involving more complex or more weakly diffracting materials. Nevertheless the combination of these factors must ensure the doubling of the data processing load on PUNCH system within the next 18 months. An important further consideration is that as the complexity of experiments increases so will the need for improved visualization techniques. Visualization is a buzz word which describes the ability to isolate the nuggets of scientific insight from the vast amount of data. In the vernacular, "to see the wood for the trees". The process tends to involve combinations of such techniques as multi-dimensional rotations with hidden line removal, edge detection and maximum entropy. These techniques are very processor-intensive and require graphics workstations to be at their most effective. It is possible that even the most powerful VAX workstation will be insufficient for some needs thus forcing the adoption of the latest ultra-fast RISC based UNIX workstations.

Conclusions

Adopting a computer strategy for the future immediately involves tackling two almost philosophical questions.The first of these is whether our users really need to store all raw data ever collected. Currently the system just copes but with the advent of large area detectors and higher neutron fluxes it will become impossible to transfer the data files by network let alone store them. Thus the answer to the question is definitely not. Instead we must concentrate powerful but flexible and reasonably user-friendly processing power at the point of data collection, that is within the DAE. This would enable much data processing particularly data volume reduction at the earliest opportunity within the PUNCH system and thus avoid the need to transfer and store large amounts of information-sparse data.

The second main philosophical question is a natural consequence of the first. If we make all of this processing power available to the FEM do we still need a HUB computer ? The answer is a most emphatic yes. Even though much raw data processing would be devolved to the FEM systems there remains a requirement for a central file storage and data administration machine. Moreover we must continue to provide a platform for a whole spectrum of more general computing to support all ISIS users irrespective of whether they are actively running an experiment. Indeed the functionality of the HUB computer should be enhanced to provide not only the greater power indicated by the empirical observation made earlier that computing requirements tend to double every three years but also more graphics facilities to aid in the interpretation of scientific results and in the production of high quality publications.The data administration role of the HUB cannot be overstressed. Currently we have about 60,000 ISIS datasets archived and this number is increasing at a rate of about 3000 per 5-week cycle. It has proven to be of great value to be able to access any one or a combination of these automatically within minutes. An example would be to compare data collected for the same material collected under different conditions and times, even from different instruments. To maintain this facility, we will require not only the HUB as a data concentration point but also a commercially supported high functionality database system such as Digital's Rdb or Compuserve's S1032. Finally the effort involved to reimplement software for a different operating system coupled with the degree of user alienation incurred makes it inconceivible that the new HUB should be any computer than a VAX running VMS.

DAE-II and Enhancements to DAE-I

Considerable thought has been given to the development of the next generation of Data Acquisition Electronics for ISIS. Much of this is described in a previous ICANS paper (Pulford,Quinton,Johnson and Norris, ICANS-X ,pp537-547). As stated earlier, the provision of processing facilities actually within the DAE would enable initial data analysis and compression without moving the data . Moreover it would provide such benefits as real time monitoring of data acquisition and detection of errors. Experience gained in running DAE-I in ISIS instruments as well as some recent technical advances have suggested some additional requirements for the intelligent DAE:

a) The DAE should be connected to the local area network, currently ethernet.This confers great operational flexibility and robustness in that DAE faults would not tend to crash the FEM as they tend to do currently. Also replacing a failed FEM would simply be a matter of using a different caller for the DAE's ethernet address. Moreover FEM's would no longer require a Q-bus and could thus be simpler and cheaper. The potential lower bandwidth available via ethernet rather than by Q-bus should be offset by the lower volume of data output by the intelligent DAE

b) The construction of the component cards within the DAE should be printed circuit rather than wirewrap as used currently. Wirewrap has proven to be very prone to mechanical faults such as broken wires or soldered pins. Also the introduction of programmable logic devices (PLD's) enables increased emphasis to be placed on producing as few different simple cards as possible, thereby improving the maintainablity of the system.

c) Predictably ISIS has been found to be the source of considerable electrical noise. Consequently long multiwire cables such as that connecting the instrument crate, into which the detectors are plugged, and the system crate are undesirable and should be replaced by multiplexed fibre optic cables.

d) As an interim measure the addressing capability of our existing MULTIBUS crates should be upgraded from 24 to 32 bits. This has proven relatively simple and has been accomplished already.

Much progress has been made in implementing this strategy. Figure 6 shows schematically two cards which have been designed and are now under construction. Their first use should be in cycle 91/1. .

Workstations

Technological innovation has brought about a rapid evolution in the capabilities of bit-mapped graphics workstations over the past few years. This coupled with considerable reductions in costs have brought a device which used to be a shared public facility for special applications to the stage where it rivals a dumb terminal in price let alone function. In recognition of this fact, a high percentage of new computer software is now written specifically for workstations. Although dumb terminals are still generally supported the future is clear, and to maintain the highest ability to analyse, visualize and present scientific results, we need to embark upon a program to replace dumb terminals with workstations on as many desks as possible.

Some of the advantages of workstations are:

1) Flexible access to more than one host computer simultaneously.

2) Potential for high quality, high speed graphics.

3) Local processing power available. The amount is generally proportional to the cost of the workstation

4) Great power and flexibility for debugging new applications.

5) Desktop publishing and word-processing.

6) All of this in one box.

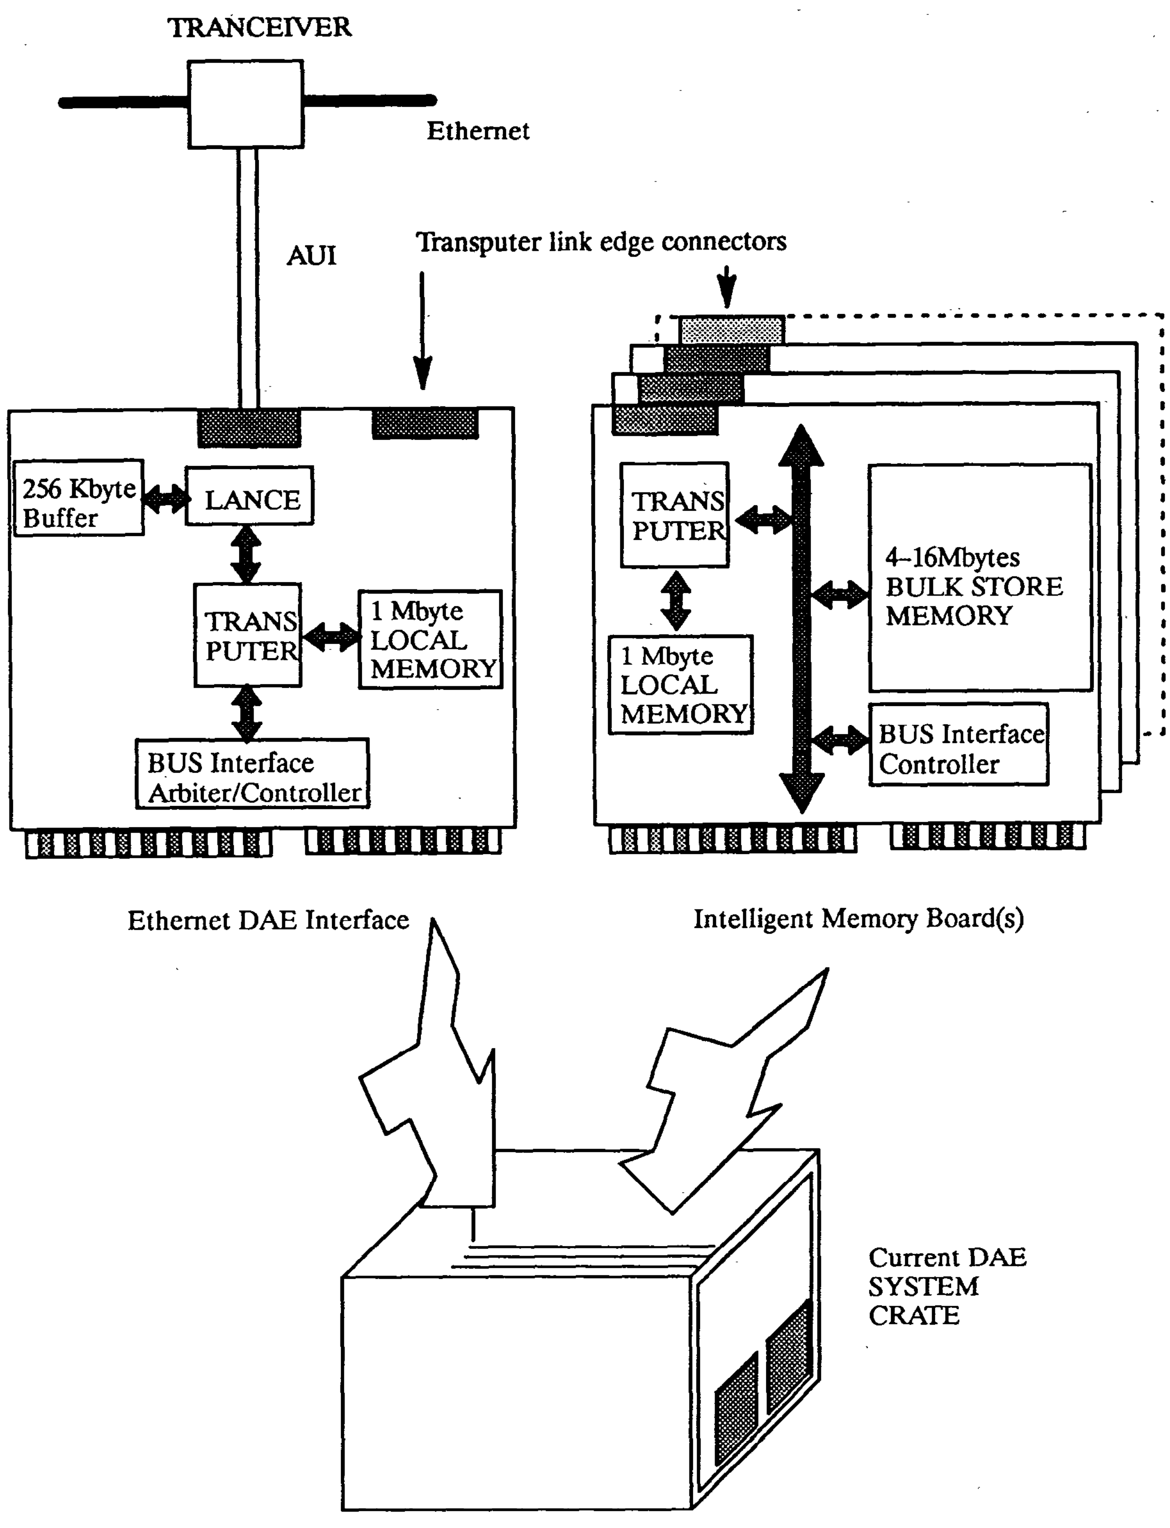

Figure 6 Soon to be available intelligent DAE cards

All of this functionality is useless unless the software controlling the workstation provides a consistent coherent interface for the user. Above all for the power of the workstation to be fully exploited, this interface must speed the interaction between the user and the system and, almost as important, be perceived by the user to do so. The same software must be common to any windowing device that we use. Currently he only viable alternative is X-windows

X windows

X-windows provides a complete working environment for a computer workstation user. Given sufficient processor power and memory it supports the creation and manipulation of any number of areas (i.e. windows) on the workstation screen including managing overlaps and overlays. Any window may occupy the full physical resolution of the workstation screen. Manipulation of the windows is performed using a pointing device such as a mouse with input coming from this device or the keyboard. The major concept behind X-windows is that processes involved with screen handling are logically divorced from the application being run. Consequently running an application breaks down into two processes: a) the X-client which is the core of the application where conventional interactive I/O statements are replaced by X-window subroutine calls ; b) the X-server which is responsible for the window display and interaction. The two communicate using a protocol called X-11 which may be run across any network link. This is illustrated schematically in figure 7a. A natural corollary is that the two processes need not necessarily run on the

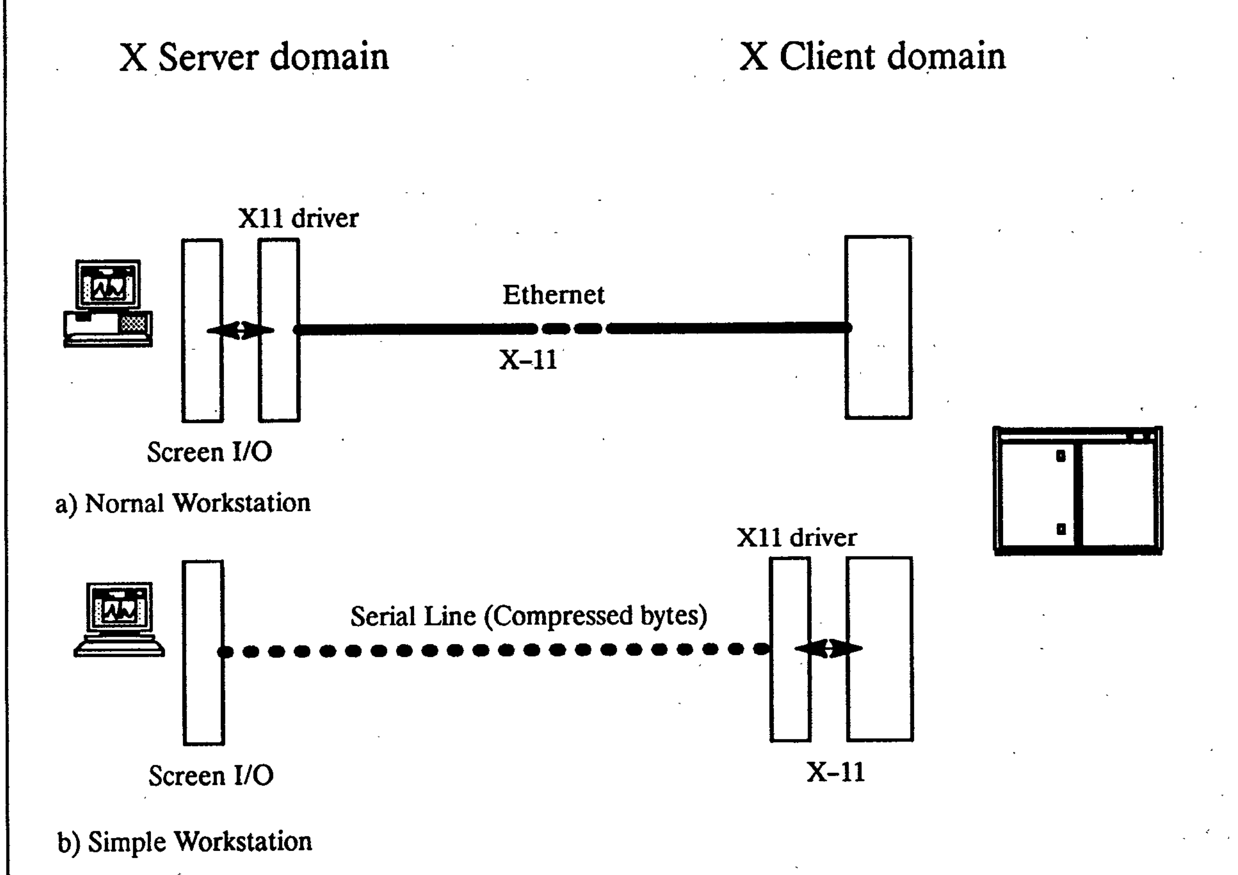

Figure 7 Schematic representation of the operation of X-windows

same computer or even one of the same type. Indeed one workstation may have a number of X-server windows each communicating with X-client processes running on a number of different remote machines. Figure 7b also illustrates a refinement to the system where the X-server software is divided into two components: an X-11 module running on the server computer which decodes the X-11 protocol into compressed bytes ready to be transmitted across a slow link such as an asynchronous serial line, and a remote windows I/O module which maps the bytes to screen operations. This scheme balances the disadvantage of imposing an increased load on the client machine against the fact that the workstation can be made cheaper to buy since it needs neither a relatively expensive ethernet port nor a large processor capable of decoding the X-11 protocol.This approach should also enable operation of windows over a telephone line. A further interesting advantage to the use of X-windows is that appropriate software (e.g. HELIOS) is available for TRANSPUTER systems. This enables a natural and integrated approach to the control of the intelligent DAE.

A Recommended Solution

The penalty for the enhanced functionality of the X-windows interface is increased demand for processor power and memory. This is illustrated in figure 8 which shows some of the relative resource requirements for two types of X-client process. The first comparison is somewhat unfair since DECwrite has far greater functionality than WPSplus however it does represent the sort of application which will be run increasingly. Figure 8b however has great significance since the running of GENIE is a fundamental requirement for any ISIS computer. The future configuration of the PUNCH system should be that which provides the most cost-effective way to supply these necessary computing resources coupled with the best value workstations.

Broadly workstations divide into three categories:

a) Simple - these provide the minimum necessary the access most X-windows applications. All desks should have one of these at least.

b) Normal - as for simple, but they have some limited processing power in their own rite. Besides running X-client software, these may have some useful dedicated applications. Some examples are IBM 386 PC running spreadsheets or MACintosh II machines running Multifinder.

c) Super - these can be expensive, but they usually provide stunning processing performance and very high resolution screens. These machines are capable of providing all resources necessary for X-windows and may themselves be used as resource providers for category a) and b) workstations.

Details of some currently available workstations are given in Table 2. The prices are approximate and are inevitably out of date. They assume monochrome monitors and no internal disk drives. The usable performance is a number based on VAX780 figures (VUP's); a somewhat arbitrary reduction factor of 0.7 has been applied to non VMS systems to account for the difficulties in exploiting the full performance of alien architecture.

Workstation	Type	Resolution	O.S.	Cost	Usable perf. (VUP)	Network
OPTIMAX	Simple	900 x 500	None	£825	none	Serial
VT1000	Simple	1024 x 864	None	£2000	none	Ethernet
IBM 386 PC	Normal	1024 x 768	MS-DOS	£2200	0.5	Ethernet
Macintosh IICX	Normal	870 x 640	Apple	£5500	1.5	Ethernet
Vaxstation 3100	Normal	1024 x 864	VMS	£6000	3	Ethernet
SUN sparc SLC	Normal	1152 x 900	UNIX	£4500	4	Ethernet
Decstation 3100	Normal	1024 x 864	ULTRIX	£6000	6	Ethernet
Decstation 5000	Super	1280 x 1024	ULTRIX	£15000	12	Ethernet

Table 2 A rough comparison of currently available workstations

The HUB must be capable of supporting a reasonable mix of workstations of types a) and b) . Any super workstation will look after themselves. An analysis of figure 8b reveals that X-GENIE requires little extra processor time but considerable extra memory and memory related resources compared with the current version of TEK-GENIE. Consequently to maintain our current level of service we need a HUB computer with perhaps slightly more processor performance but with greater physical memory and large capacity fast disk drives for paging. DEC have recently announced the VAX4000, an ideal machine for this purpose which is particularly optimized for distributed applications such as X-windows. Table 3 gives a brief summary of the major parameters of this system. Clearly in

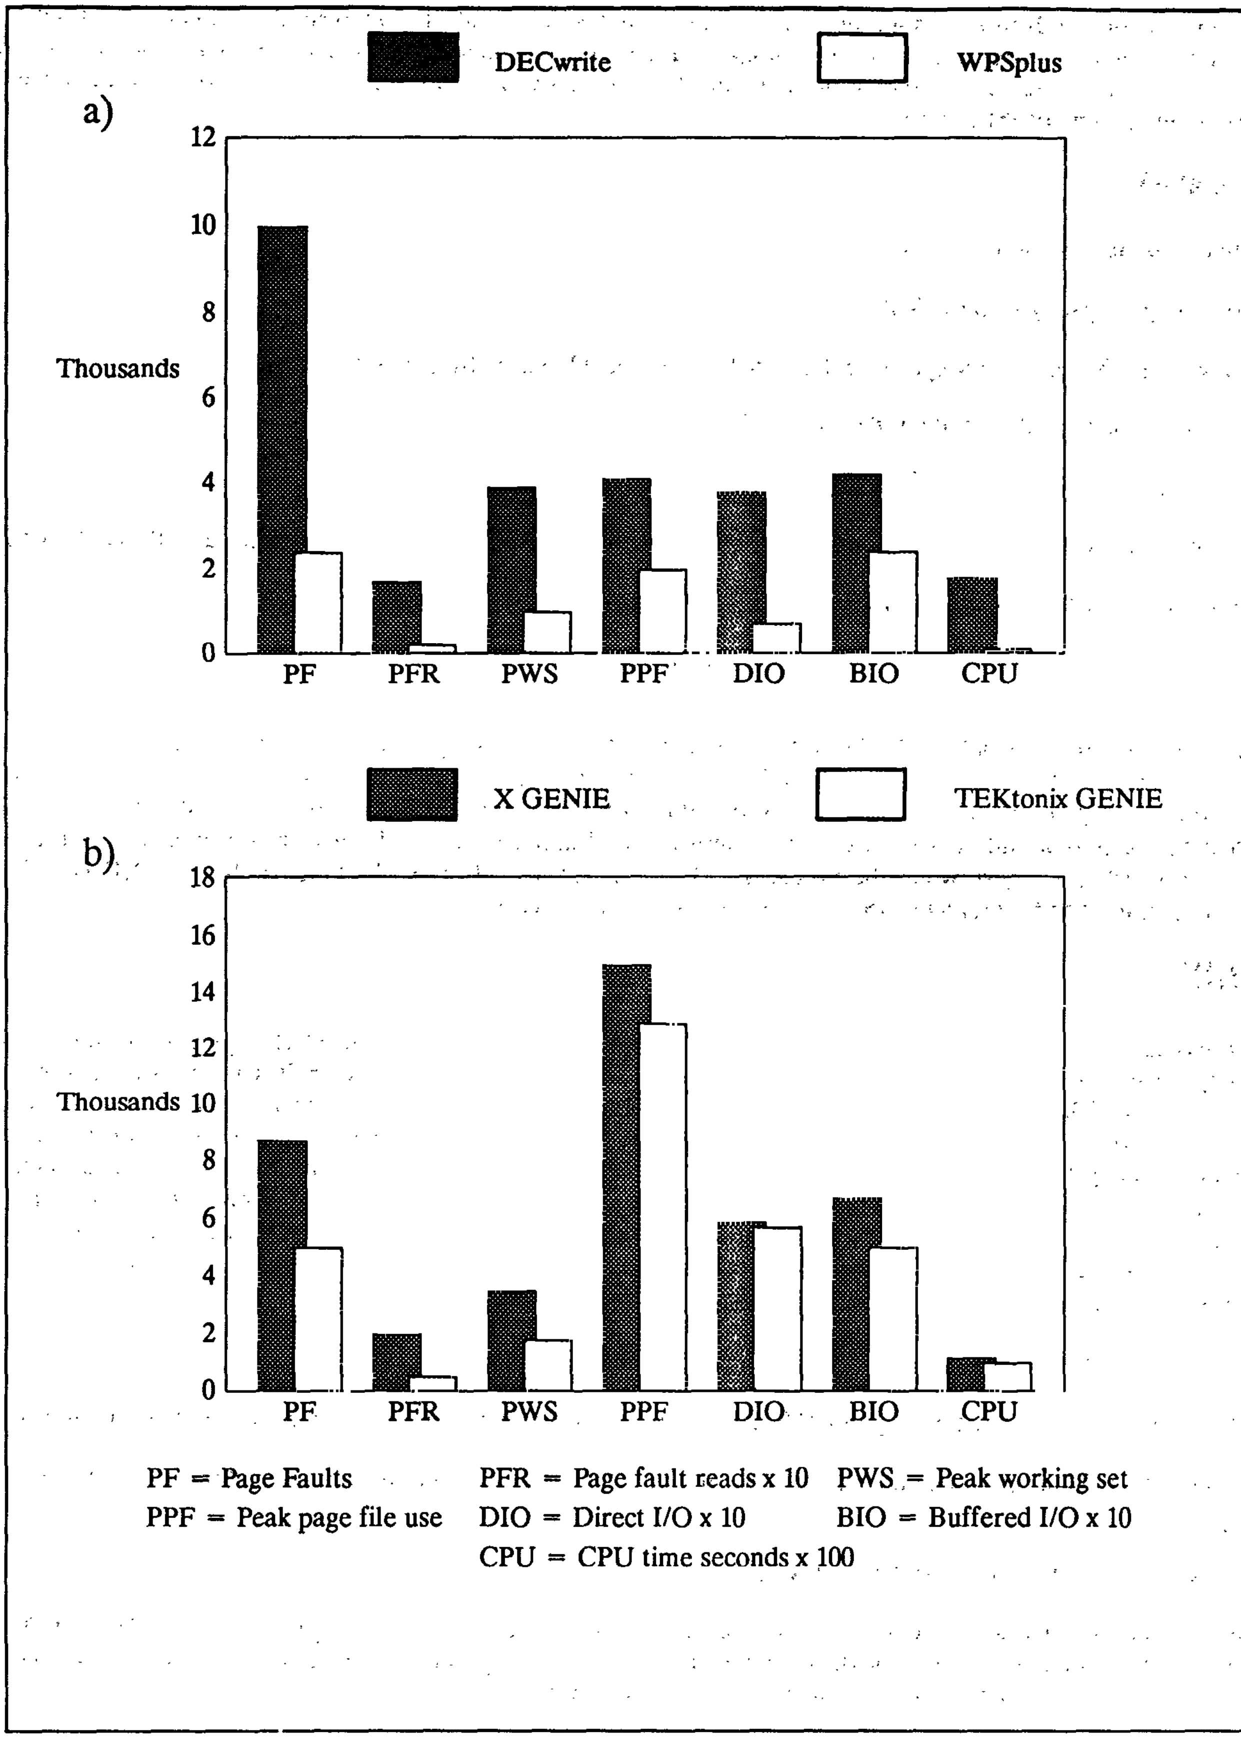

Figure 8 A comparison of the key system parameters with and without X-Windows two applications on a simple workstation

anticipation of our future computing requirements, this system should be duplicated to double the available resources. Judicious choice of disk technology would enable the possibility of disk sharing between the two processors. This in turn leads to considerable fault tolerance, a commodity which has been found to be of value during the running of ISIS up to now.

VAX 4000

Memory 32Mbytes -> 128Mbytes

Disk storage capacity - up to 28Gbyte

Bandwidths - Memory 40 Mbytes/sec, DSSI 8.0Mbytes/sec, Q-bus 3.3Mbytes/sec

CPU performance - 8 x VUP (cf. 8650 6.5VUP)

Power consumption - 0.9kw. (cf. 8650 8kw)

Cost of 1 x VAX 4000 and 5 Gbyte magnetic disk - about £100,000 as as interactive machine, £80000 as SERVER.

Maintenance - about £4500 per annum (cf 8650 £25000 per annum)

Table 3 Introducing the VAX 4000

Figure 9 illustrates a possible PUNCH configuration of the future. The diagram also shows an enhanced central data highway based on the 100Mbits/second FDDI fibre optic token ring. This, although not needed immediately, would be required to provide the necessary bandwidth for future data transfer. The computer configurations within instrument cabins would be associated into groups to avoid unnecessary local network traffic escaping onto the network backbone. These would be probably : a) MUSR,LAD,TEB,MARI b) POLARIS,EVS,SXD,HRPD c) SANDALS,PRISMA,LOQ,ROTAX and d) IRIS,TFXA,HET,CRISP.

UNIX

There is no doubt that the most powerful computers for a given cost are based on the UNIX operating system. As an example, we have found that the ULTRIX based DECstation 5000 gives twice the VAX 8650 processor performance for a portable FORTRAN application. Subjectively UNIX workstations like this also provide very rapid response for windows servers. Unfortunately the UNIX operating system, although much loved by computer scientists and very elegant in many concepts,, cannot be said to be user-friendly and it would be unrealistic to subject our users to it without copious health warnings. The editors are usually powerful and dangerous. Table 4 provides a few example commands and their VMS counterparts to emphasize their abstruse nature. Nevertheless we should incorporate some of these devices into the future PUNCH system to exploit their great processor power. The inclusion of a UNIX machine must be subject to strict criteria:

1) It must run X-windows..

2) Its native compilers, particularly FORTRAN, should compile standard VMS source code without modification

3) It must be possible to submit a job from VMS to the UNIX machine and retrieve the output in a well-defined, simple way.

4) Its speed should make 1,2 and 3 worthwhile.

DEC themselves have acknowledged the potential of UNIX(ULTRIX) computers in producing the VMS-ULTRIX connection software. This enables VMS and UNIX to share files transparently across a network as well providing for shared printer and batch queues. We should make full use of this.

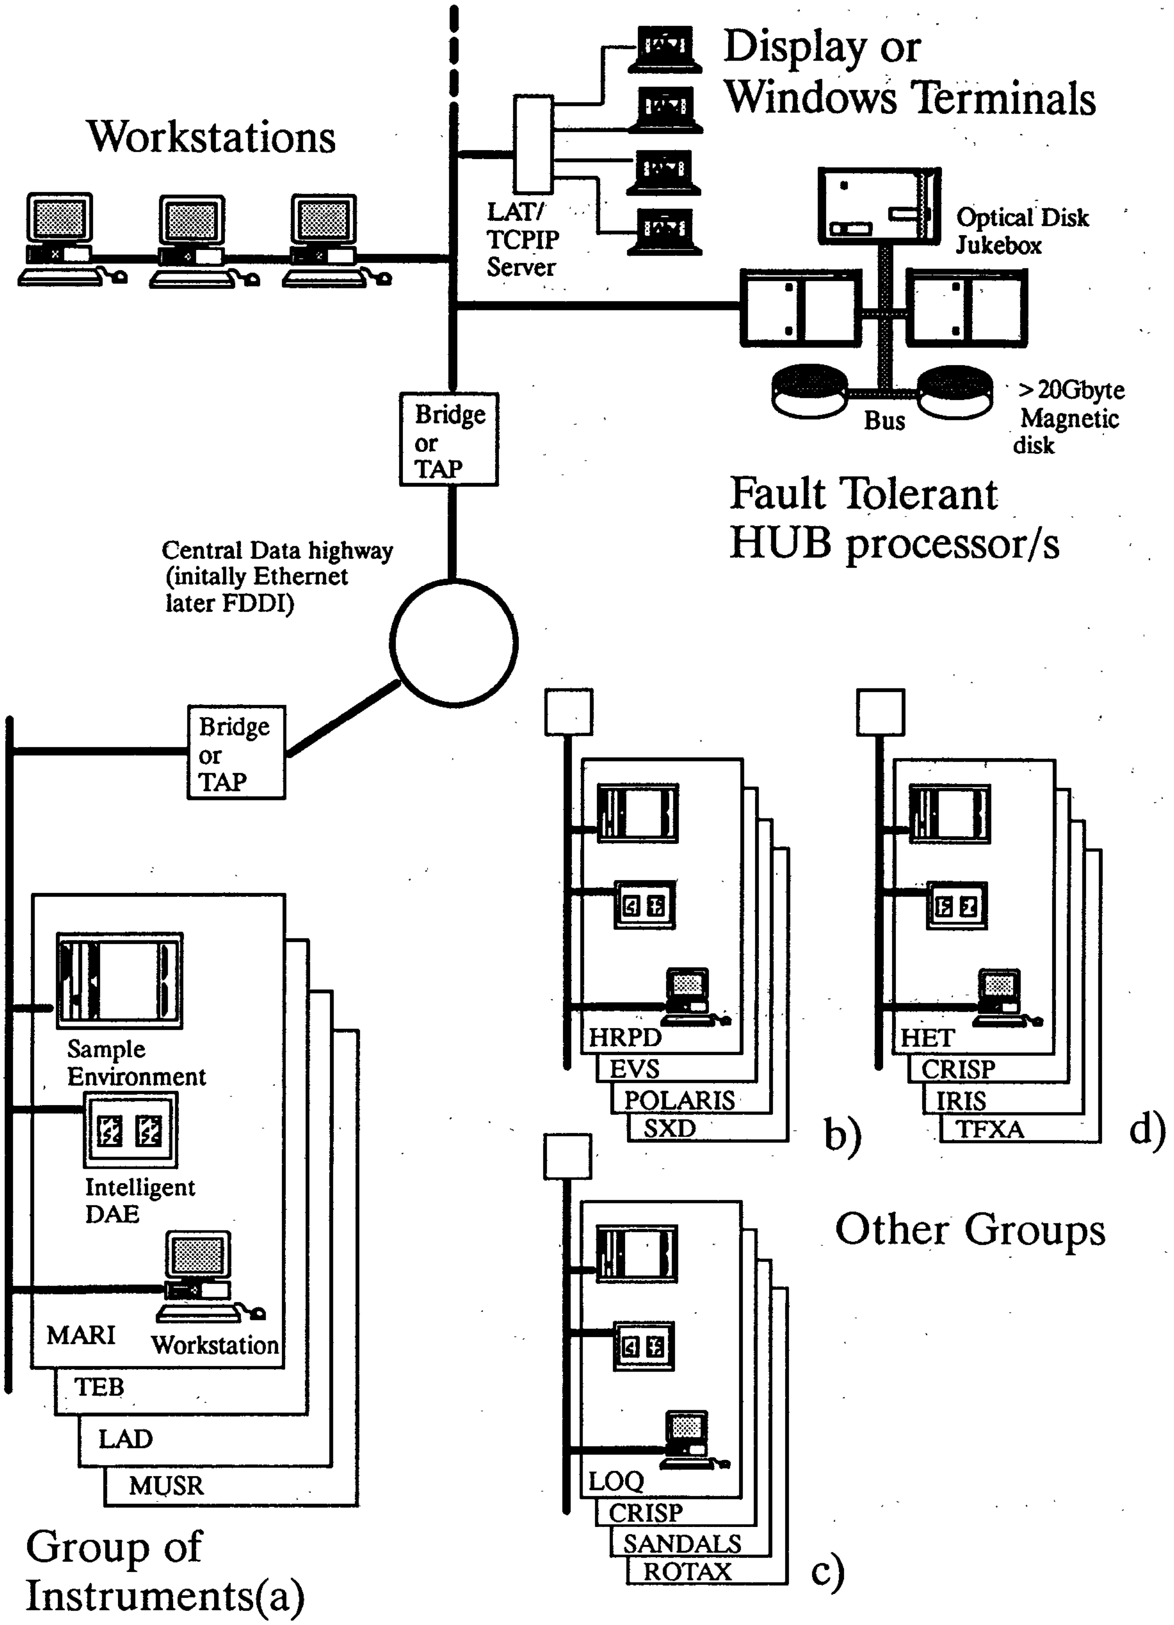

Figure 9 Probable future configuration of PUNCH

UNIX filename *file* - pathname/filename

VMS - DEV:[directory]filename

VMS COMMAND	UNIX COMMAND
$ APPEND *file1 file2*	$ cat *file1* > > *file2*
$ COPY *file1 file2*	$ cp *file1 file2*
$ CREATE/DIR *[dirname]*	$ mkdir *dirname*
$ DELETE *file*	$ rm *file*
$ DELETE/CONFIRM *file*	$ rm -i *file*
$ DIRECTORY	$ ls -a
$ DIR *file*	$ ls *file*
$ DIR/PROTECTION/DATE/SIZE	$ ls -al
$ EDIT *file*	$ ed,ex or vi file (to introduce a powerful unforgiving editor)
$ MAIL	$ mail
$ PRINT *file*	$ print (or lpr or pr) *file*
$ SHOW DEFAULT	$ pwd
$ SHOW PROCESS	$ printenv
$ SHOW TERMINAL	$ stty everything
$ TYPE *file*	$ cat *file*

Table 4 A comparison of some equivalent UNIX and VMS commands

A Future Workstation Screen

The final section of the paper is mainly to provoke thought and discussion. Within 1 year it will be possible given sufficient skilled manpower and resources to produce the underlying applications to format a workstation screen such as that shown in figure 10. The following list of features might be supported:

1) Operations such as changing a sample environment parameter value would be achieved by pointing at the relevant box, clicking a mouse button, and typing the new value on the keyboard.

2) The client applications would run on any one of a number of different hosts. As an example, the real-time picture of the detector would be the output of a process running on the intelligent DAE.

3) Producing wiring tables and any such gemetrical applications would be accomplished by simply pointing using a mouse.

When this notion was presented at a recent meeting of ISIS scientists it provoked a great diversity of response such as:

a) "I'm confused enough as it is!"

b) "It would make the control of instruments much easier for the users"

c) " Just typing a three-letter mnemonic in reply to a terminal prompt is far more efficient"

d) " My arm will get tired!"

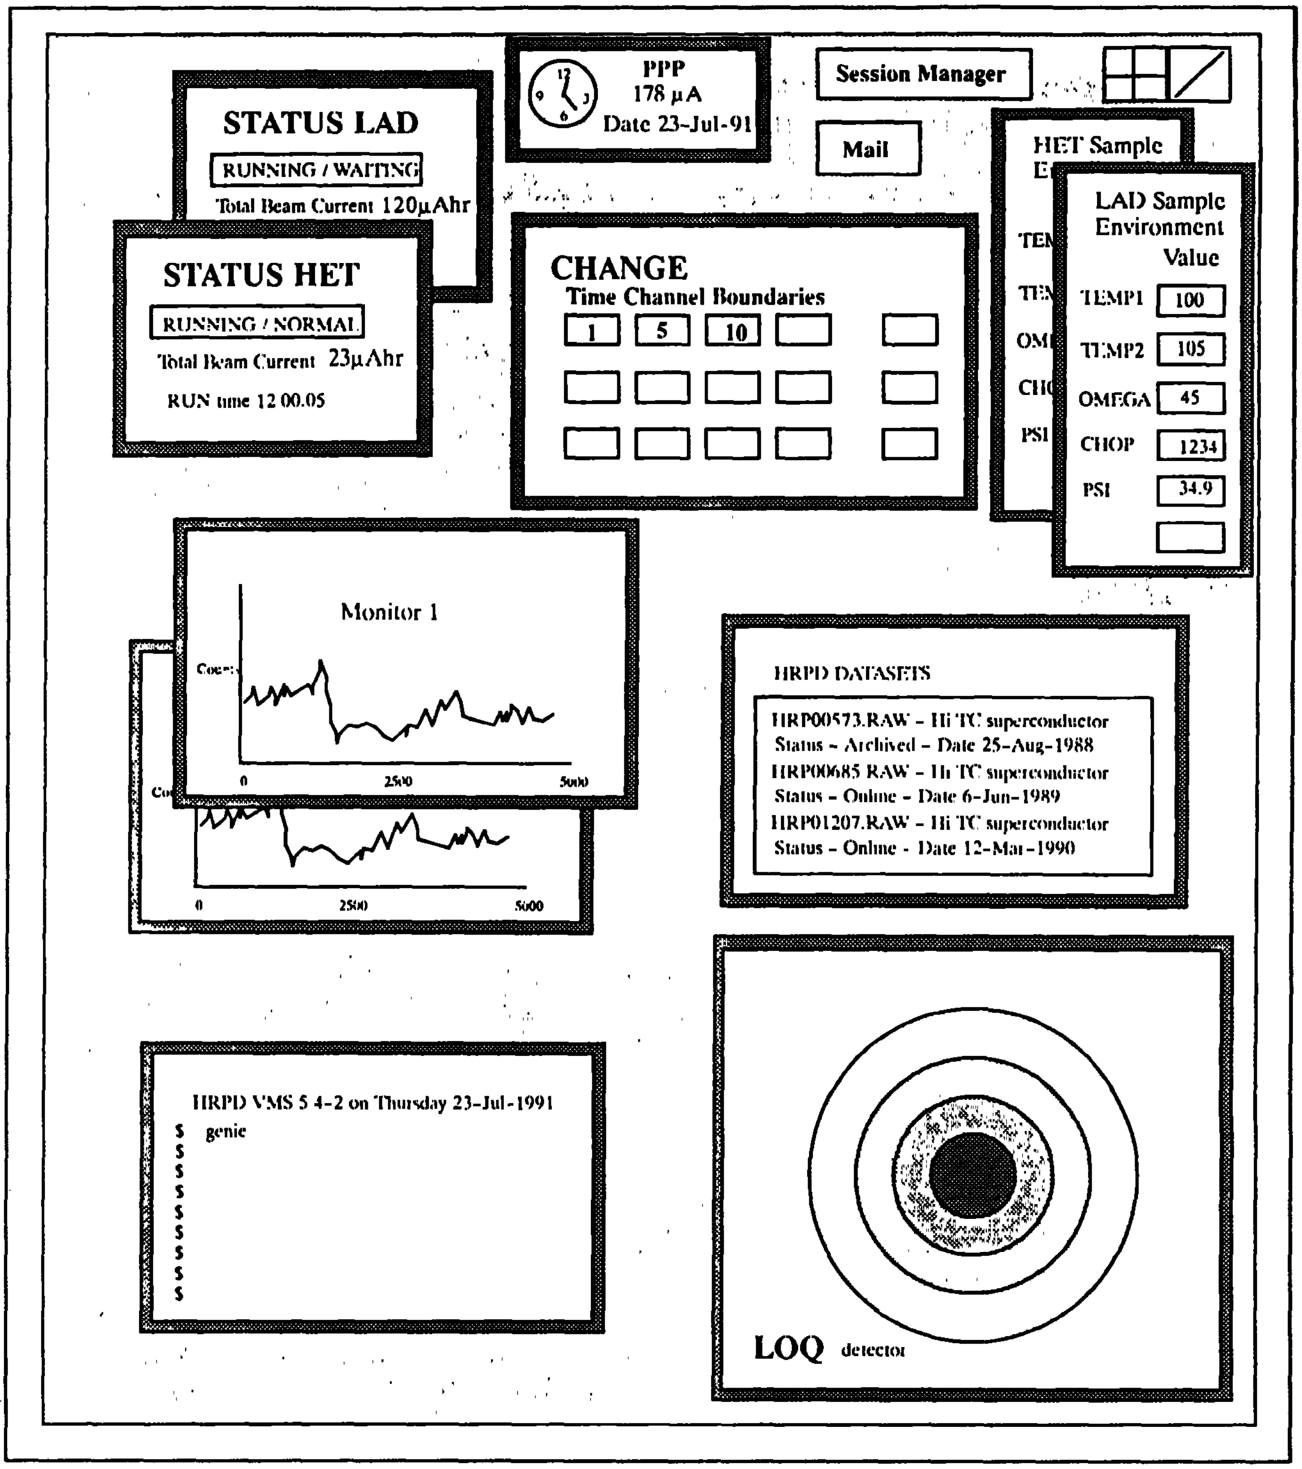

STATUS LAD
RUNNING / WAITING
Total Beam Current 120µAhr
STATUS HET
RUNNING / NORMAL
Total Beam Current 23µAhr
RUN time 12 00.05
PPP
178 µA
Date 23-Jul-91
Session Manager
Mail
HET Sample
LAD Sample Environment
Value
TEMP1 100
TEMP2 105
OMEGA 45
CHOP 1234
PSI 34.9
CHANGE
Time Channel Boundaries
1 5 10
Monitor 1
0 2500 5000
HRPD DATASETS
HRP00573.RAW - Hi TC superconductor
Status - Archived - Date 25-Aug-1988
HRP00685.RAW - Hi TC superconductor
Status - Online - Date 6-Jun-1989
HRP01207.RAW - Hi TC superconductor
Status - Online - Date 12-Mar-1990
HRPD VMS 5 4-2 on Thursday 23-Jul-1991
$ genie
LOQ detector

Figure 10