

MCD Compute Resource Strategy

Anthony Scott-Hodgetts

January 22, 1988

It would be very gratifying to spend a week exploring possibilities for MCD Compute Resource, to write a concise report on my conclusions, and to have MCD act upon those recommendations. However, this would be pointless. We already have a perfectly acceptable strategy document, authored by Clive Dyson. Sadly, his report which was aimed at the wider issue of Inmos Compute Resource was never acted upon for lack of funds. I am resubmitting this document as a basis for the future MCD Compute Resource Strategy. I do not believe in re-inventing wheels!

The major points for discussion are: which networking protocol is most appropriate, and what machines are we forced to buy in order to support specific pieces of software, for example, externally written CAD packages.

1 Introduction - current problems

As a company we are currently suffering from the following computer related problems:

1. Insufficient computing power to allow necessary tasks to be performed (especially in manufacturing support). This is resulting in inefficient use of our primary resource - good engineers. A side effect of this is that we may be losing some of our best staff due to their frustration at not being provided with the tools they need to do their job.
2. Isolation of data on incompatible makes of computer. We have no strategy in place to ensure that all data in the company can be accessed by the people needing to refer to it.
3. Lack of enforced mechanisms to ensure that all data is backed-up and those copies are stored at remote locations. This costs time and effort when disks which have not been backed-up crash.
4. We are not achieving the best discounts available to us because Individual groups are not encouraged to coordinate their hardware and software purchases. This also results in us constantly duplicating effort locating suppliers for commonly used items, and in finding out how to install and use them.

5. Development of different working environments and practices. We use different editors, user interfaces and other software tools on different sites. Most sites use a range of all of these. This means that our files are generally incompatible, and that staff moving between sites have a bewildering range of different environments to learn. Also, the task of software maintenance is effectively multiplied by the number of sites. (Alternatively this can be viewed as reducing the efficiency of our systems programmers by the number of sites.)
6. Divergence of hardware for different groups. A good example of this is the use of different hardware in the CAD and software support groups in MCD. This makes the development of common software difficult, although we are trying to establish a closer working relationship.

These problems need to be addressed by a coordinated plan for future computing development which is directly geared to staff movements and overall levels of capital expenditure. The following is a proposal for such a scheme. It puts forward a set of strategic aims, describes a possible implementation and outlines the resulting benefits to the company. This is followed by a costing for the resources required, leading to a very approximate yearly capital budget for the company.

2 Strategic aims - part I

2.1 A hierarchy of filestores

As a general strategy we should aim to provide a hierarchy of filestores to hold the databases within the company. (The term 'database' in this document refers to a collection of data which must be viewed as a single entity, rather than to a database program.)

Consider, for example, our UK manufacturing database. This is very large, and has to be consulted by a wide range of people in Newport and Bristol. Such a database is currently best served by a large, single, filestore running on a large VAX with high input/output capability (e.g. an 8350). However, as will be noted in the next section, it is important that CPU intensive tasks such as simulation are not run on the same machine.

Other databases are not as large, and do not need to be consulted by the same range of users. For example, the database for any of our chip designs is of the order of 100 to 500 Mbytes, and need only be consulted routinely by a team of fewer than 10 engineers. It can easily be held on a small machine (such as a MicroVAX) with the appropriate disk capacity.

Many other 'departmental' databases within the company fall into this category, and could easily be supported by a local filestore.

Finally we come to the collection of files usually held by a single user. It is tempting to follow the above hierarchy to its logical conclusion and provide each user with a local disk for their own data. However, this must be avoided at all costs. This is because it is notoriously difficult to ensure that such disks are adequately backed-up, and they tend to prevent the sharing of data between members of working teams which helps to prevent the duplication of effort and facilitates communication. Thus the lowest level of filestore we should aim to support is the local group filestore, which should hold the database for a team of from 10 to 40 users, depending on requirements.

2.2 A hierarchy of processing resource

As mentioned above, it is generally not a good idea to overload large VAXes holding large databases with other CPU intensive tasks. The costs of raw MIPS is falling so rapidly that we should aim to put the processing power as close to the user as possible. This trend is already clearly visible within the computing industry, and we are of course intending that the transputer is part of it.

Thus all users processing requirements should be removed from our central VAXes, except when the processing is related to the manipulation of large databases resident on those VAXes.

The manipulation of the data stored on local departmental filestores can be carried out on the filestore host, whilst highly CPU intensive tasks, such as graphics, simulation, spreadsheets, and even word processing, can be performed on a workstation allocated to each user.

Clearly, casual computer users (this does not include secretaries) do not need their own workstation. They can be supported by a dumb terminal connected to a local host, or a central VAX as appropriate.

Thus our strategy should be to distribute our processing power in an inverse hierarchy to that used for our filestores. That is, provide the processing power required as close to the user as possible, consistent with the location, and the quantity, of data being manipulated.

2.3 A consistent company-wide communications network

If users are to be distributed on a hierarchy of computing resource and filestores, as described in the previous two sections, then we must provide a single communications network to allow all users to access all computers

in the network. This will also allow companywide communication between users.

Fortunately, DECNet provides a readymade basis for such a communications network. It can handle all our inter-site communications, as it does now. All that we need to ensure is that any filestores or workstations that we connect to the network are compatible with its protocols.

Networking software to allow computers from manufacturers other than DEC to be connected to the network via Ethernet is now available. We are already connecting IBM PC's to our VAXes by DECnet DOS, and we should be able to use NFS (Network File System) and DNA software to connect in Sun workstations as well. Note that this communications network will allow IBM PCs to logon to both VAXes and SUNs and gain access to all tools on these machines.

Our strategy should be that all users have direct access to the network. It is therefore a requirement that any workstation purchased can be connected to, and used with, our company communications network.

3 A practical computing hierarchy

The three strategic aims presented in the last section can be satisfied by extending our current computing resource incrementally towards the computing hierarchy described in this section. Note that implementing the proposal in an incremental manner means that we can spread the capital cost over a number of financial years (say 3), rather than having to make a single one-off purchase of replacement equipment.

At the top of the computing hierarchy are our current VAX clusters, one on each site. These should be used to hold our large databases, and new large VAXes should only be purchased for this purpose.

Inter-site communication is provided by DECnet. The current inter-site links that we use are relatively low bandwidth, and the upgrade of these links to 1 Mbit/sec bandwidth should be given a high priority.

The next level in the hierarchy is a number of departmental hosts and filestores (see Figure 1). These will typically be MicroVAXes with the disk capacity needed to support a given departments' data storage requirements. (In the future it may be possible to use other hosts, such as Suns, if this proves to be more cost effective.) All of these local hosts are connected together via a central, site-wide, Ethernet, which is also connected to the sites' duster of large VAXes.

Each filestore will host a group of workstations to support the users within a department. Each group of workstations is termed a 'local network', and IBM PCs, Sun workstation, VAX workstations, and CAD workstations can be connected to it as required. Each local network provides its users with the majority of the computing resource and file storage that they require, and also provides access, via the site-wide Ethernet, to large databases stored on the central VAX dusters. Access to other central resources, such as large plotters and tape drives, can be obtained in the same way.

Note that a dumb terminal (e.g. a VT220) can be connected directly to the local host via a serial line rather than via the local network.

Figure 2 shows a local network in more detail. The majority of local networks in the company (those that do not use transputers) will use thin Ethernet to provide their local communications. We will be able to provide transputers within the MicroVAX local host, and within the IBM PCs and Sun workstations, and these transputers will also be able to communicate via a link-based network ("TransNet").

CAD workstations will be purely transputer based, and will communicate with the host via TransNet. We will have the option of providing silicon designers with a Sun workstation to act as a terminal to the CAD workstation, if they require using external vendor's software.

4 Strategic aims - part II

Having described the computing hierarchy that we propose, this section considers some further strategic aims we would like to satisfy.

4.1 Future-proof hardware

Note that the hierarchy described above is dependent upon a consistent communications network, and common file interchange protocols. It also requires users to be able to gain access to log on to a VAX so as to use the facilities of DECnet for electronic mail etc.

However, it places relatively few constraints upon the hardware provided to act as a user's workstation. We should therefore be able to change our favoured hardware as the technology improves, which leads to:

4.2 A 'Common Base' policy

At present managers are free to choose the hardware they purchase. This leads to a proliferation of different computers, which cannot be easily connected, require a large duplication of effort to support, and are expensive to maintain (as we cannot negotiate good maintenance contracts or carry even a small range of spares).

It is proposed that we establish a 'Common Base' policy. This would involve a representative working group establishing a range of preferred hardware. Purchase requisitions for preferred hardware would be approved with far less red-tape than for non-preferred hardware. As we would have a good idea of the quantities required in any year we should be able to negotiate extremely good terms for both purchase and maintenance.

Consider as an example the use of IBM PC compatibles within the company. We should specify one make (and model) of desktop PC AT compatible as preferred and one make and model of a portable compatible as preferred. These need not be the cheapest on the market, as the saving available on purchase and maintenance will allow us to choose the one most suitable for our company wise needs.

Note that such a common-base policy is equally applicable to software, and that we may be able to obtain substantial discounts on site-wide licenses for software such as spread sheet calculators.

4.3 Classification of users

The majority of INMOS employees have well defined requirements for computer resources. By defining the hardware which may be allocated to an employee we can ensure that the required hardware can be easily planned and budgeted for. Furthermore, whenever any new member of staff joins, the procedure for providing them with the appropriate computing resource is automatic.

The decision as to the hardware a user requires can be left to their local manager, who can decide on the appropriate equipment from that specified in the Common Base. Assistance could be provided to local managers in this task from a central advisory group.

A possible set of hardware is given in the following table:

Users Hardware	Typical users
Dumb terminal	Casual users making little use of computing resource, engineers who only consult large databases which expect a terminal interface, in the short term many users in the next class.
IBM PC	PCB design engineer, secretaries, directors, senior managers, project leaders, FAE, PME, test engineer, product engineer, and QA engineer.
Sun workstation	Software engineer, technical author.
CAD workstation	Silicon design engineer.

5 Benefits of the above strategy

1. CPU intensive tasks which prevent the efficient operation of the large databases on our large VAXes are removed and located on relatively cheap processing power close to the user.
2. Everyone in the company has access to supported facilities including: electronic mail, word processing, central databases, automated back-up, retrieval of data from archive, and management tools.
3. Disk space is minimised for a given amount of data, as shared data need only be stored once.
4. Costs are minimised by bulk purchasing and reduced maintenance charges.
5. Our engineering resources are used more efficiently, by ensuring that tasks are only performed once, and common hardware and tools are well supported.
6. Users are protected against loss of data through supported back-up.
7. Software groups within the company are encouraged to coordinate their work.
8. Finally, a number of less quantifiable, but important benefits such as improved efficiency, improved service to customers, improved morale, and better quality.

6 How much will it cost?

The strategy described in this document makes budgeting for our computing needs a relatively simple task. We can separate the budget into two sections:

that required providing central VAX clusters to support our large databases; and that required providing local networks for individual departments within the company.

For the present this document ignores the costing of the VAX clusters and concentrates on the distributed computing requirement.

In this area we are starting from a relatively small installed base, and for simplicity this is taken to be zero to make the sums easy.

We currently have a worldwide staff of around 1000. Assume that we need around 30 local networks to support this workforce, with 15 workstations per local network (I need to talk to personnel to check these assumptions).

Each local host and filestore costs around £60k, and each workstation costs an average of around £6k. Thus our total final companywide requirement is for an installed capital base of around £4.5M. Clearly this is a very approximate estimate, but the method of calculation is very dear.

It would be impractical, and risky, to install this capital base in a single financial year. It would be preferable to phase the installation over three years, spending £1M in the current financial year, £2.5M in the second year, and £1M in the third year. Expenditure in subsequent years should run at around a quarter of the installed capital base, plus an additional sum to cover any additional staff joining the company.