ALVEY DIRECTORATE INFRASTRUCTURE POLICY

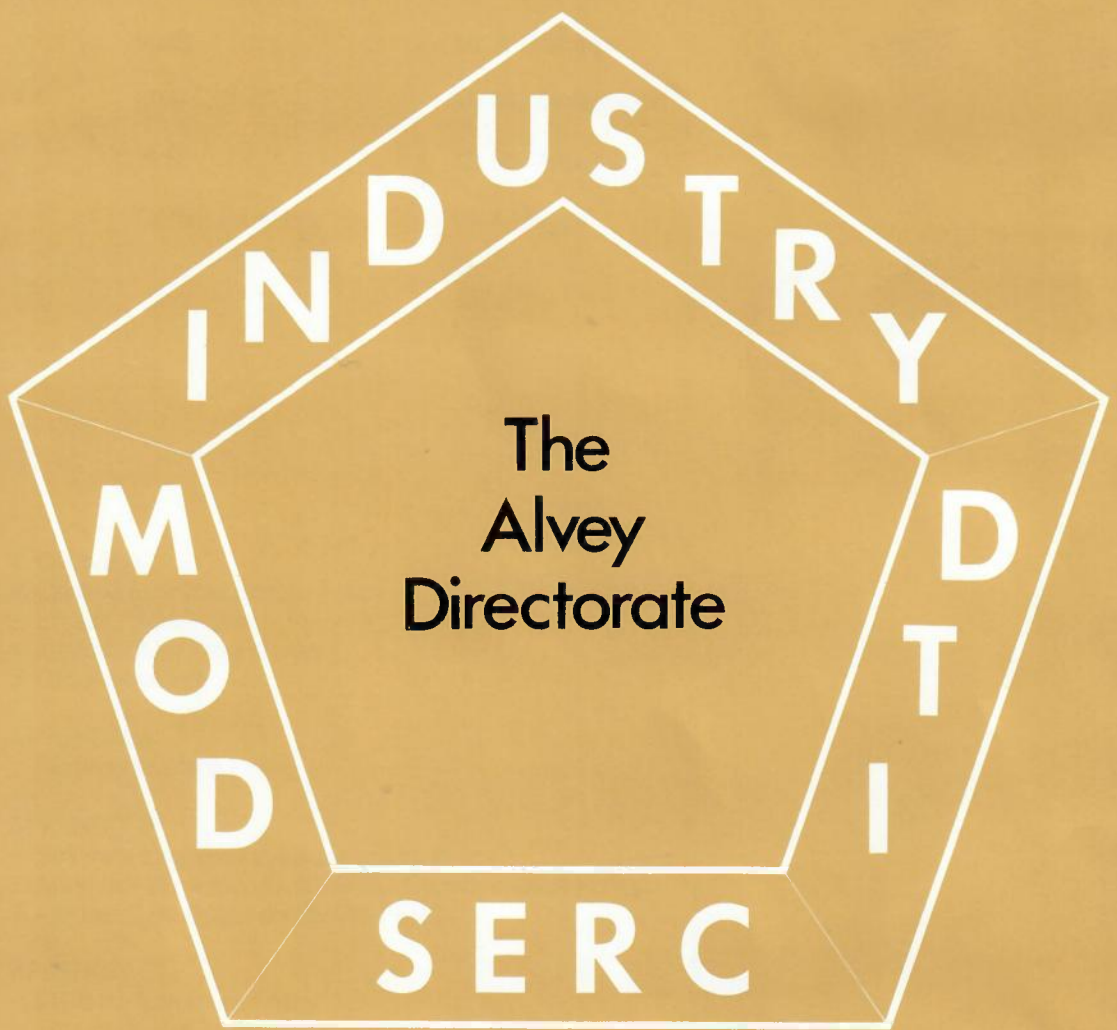

IEE
Published by
The Institution of Electrical Engineers

Alvey Directorate
Dept of Trade and Industry
Millbank Tower
Millbank
London SW1P 4QU

September 1984

PREFACE

The Alvey Directorate intends to coordinate its provision of computing resources for projects by setting up an Alvey Infrastructure Programme. This report is the output from a small working party set up within the Directorate to produce the initial plans. It is now being issued as a consultative document so that the views of participants in the Alvey programme, current or prospective, can be obtained during the detailed planning stage.

Comments on the main body of the report should be sent in the first instance to W P Sharpe, Alvey Directorate. Comments on the appendices should be sent to the representative of the area concerned:

Software Engineering	Dr R W Witty
IKBS	Mr W P Sharpe
MMI	Mr M J Underwood
VLSI	Mr M J Newman
Communications	Mr A Griggs

CONTENTS

1. Overview

This report sets out the policy that will be followed by the Alvey Directorate in providing the hardware and software computing infrastructure upon which the work of the programme is carried out. The purpose of the policy is to bring together the requirements of the various parts of the Programme and to define a common framework of standards and supported services that will enable effective collaboration and generally improve productivity.

Section 2 of the report establishes the context for the policy by considering the place of the Alvey Programme within the wider industry infrastructure. The main classes of users and the components of infrastructure are defined and the focus of Alvey policy is identified.

A collaborative programme is only possible if certain standards are laid down for the languages, tools, etc. The Directorate is not a standards organisation so in this context the term refers not only to formally ratified standards but to those guidelines and specifications that may be adopted for use within the programme. Standards enable and promote the interchange of ideas as well as artefacts. They also serve to reduce the waste of effort that results from needless duplication of development work. The Alvey Programme must be concerned with both the productivity of its directly funded activities and with the effectiveness of exploitation. To meet these requirements section 3 lays down the general guidelines of the policy on standards and choice of equipment on Alvey projects; detailed statements of standards for each Alvey category are contained in the appendices.

The standards policy is the necessary basis for a programme of coordinated investment and support that the Directorate intends to establish to meet the computing requirements of individual projects. An overview of this programme is given in section 4 and detailed category plans are in the appendices.

Finally, section 5 covers the interaction of the Alvey infrastructure policy with various other projects such as AWSAP, SERC Common Base, etc.

Some parts of the Alvey programme, such as Software Engineering, were preceded by a long period of planning and have been able to build on an existing statement of policy. Others, such as MMI, are still involved in the definition of their programme. The report therefore cannot be complete at this time and it will be reviewed regularly as plans develop and needs change.

2. Context for Infrastructure Policy

In setting out a computing infrastructure policy it is useful to distinguish three classes of user and the types of system they use:

(a) The vendors produce, and the users of IT products require, delivery (target) systems.

(b) Research and development in IT utilises development (host systems).

(c) Research into the virtual machines (hardware and software) that support applications will frequently require prototype development and delivery systems, i.e. the bulk of Alvey R&D work will be done on systems providing a full, stable, development environment, not on systems that have been prototyped in order to allow investigation of some new architectural feature (e.g. data flow).

A further categorisation of the main components of the infrastructure into systems, communications and software allows us to draw the following diagram.

The boundaries between the various types of system are not hard and fast but may be roughly characterised as follows.

Single User System (SUS) are distinguished by having special resources devoted to supporting user interaction, e.g. bit map graphics, touch screen, etc. By attaching these resources to a dedicated personal computer, systems are built that currently lead the way in facilities for interaction and hence for the productive use of human resources. The ICL PERQ is a typical example of this type of system.

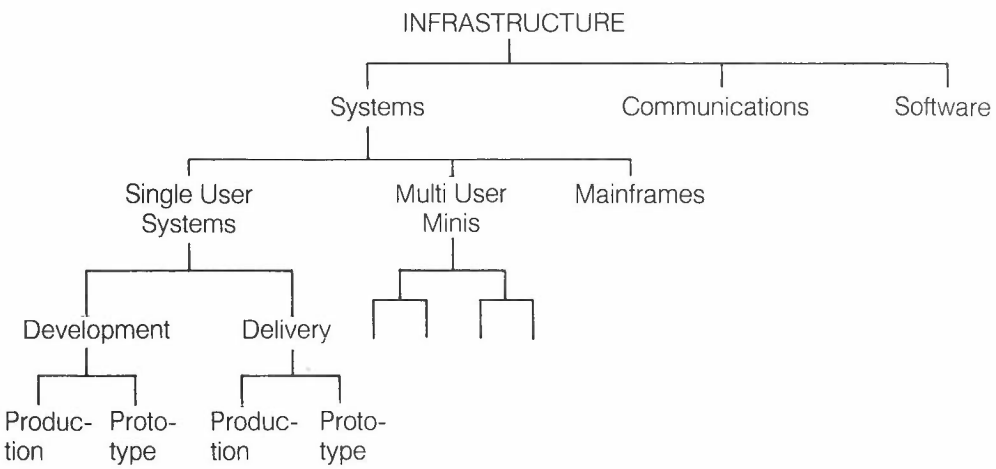

Multi User Minis (MUMs) are generally able to deliver more power (CPU speed, memory, backing store, etc.) to any single problem than can the SUS, but are not able to match their interactive capabilities. A MUM at the lower end of the performance scale might be no more powerful than a SUS; such a machine is not usefully regarded as a SUS in the context of this policy.

Mainframes represent the top end of the multi user performance scale.

The Alvey Programme is concerned with pre-competitive research and development in the enabling technologies, and with bridging the gap between R&D and exploitation. The emphasis of Alvey policy and activity for the three classes of user identified above will therefore be as follows:

(a) The choice of Delivery systems is a concern of the vendor. The Alvey Programme will be concerned with the prototyping of innovative Delivery vehicles but not their production. Normal DTI funding mechanisms will continue to apply for taking prototypes into production. The users in this class require that their investment in application software is protected and they will benefit from the software standards policy insofar as its influence extends beyond directly funded Alvey R&D into the production and exploitation of Delivery systems.

(b) The main part of this policy statement is addressed to the needs of this class. Research in AIT tends to go hand in hand with development of the programming languages and environments in which programs are developed and expressed. It is therefore necessary to weigh the benefits of a common research infrastructure against the disadvantages of over rigidity. A flexible policy that supports the standards, reduces needless incompatibility, and brings new advances into the infrastructure is necessary.

(c) The prototyping of new development or delivery machines is an important element of the Alvey Programme but it must be distinguished from providing Production Development machines. Insofar as the users in this category make use of a common research infrastructure their needs are identical to those of class B.

3. Infrastructure Policy

3.1 Standards Policy

The standards policy is concerned with both software and hardware. Hardware standards relate to bus standards (e.g. IEEE 488), etc., rather than complete systems; statements on complete systems are made under systems policy.

The purpose of a standards policy is twofold:

With respect to R&D it should protect investment in tools so that new hardware developments can be taken advantage of with as little disruption as possible, and it should aim to maximise cooperation between all areas of R&D.

With respect to exploitation it should smooth the path from R&D prototype to delivery systems and protect investment in applications packages.

Statements of Alvey standards will fall into the following categories:

(1) Mandatory Standards

A specification which is sufficiently stable and widely available for it to be used for all stages of Alvey funded R&D. These specifications will usually have formal standards status or be the subject of Alvey supported standards activity.

There will also be proprietory standards in this category which it is recognised will not have formal standards status (e.g. NAG library subroutine interface).

(2) Preferred Standards

A preferred Standard is a specification that the Directorate would like to see used as widely as possible within the Programme but which cannot be given mandatory status in the short term. The weaker status of a Preferred Standard is a recognition that it is either not a formal standard, or is not sufficiently widely available for there to be any practical purpose in declaring it to be mandatory. Adherence to the standard is required wherever it is a realistic option.

Migration from preferred to mandatory status will be normal and the infrastructure policy will be regularly reviewed to identify areas where a stronger standards activity would be beneficial. The designation of a specification as having Alvey standard status means that there is a commitment to it from the Directorate which will be given practical expression by:

(a) Direct investment to make it available to the R&D community on Supported systems,

(b) Support for conversion of tools to the standards on Supported systems,

(c) Support for conversion to a new standard if the current one is superseded,

(d) Encouragement for its support through the other IT funding mechanisms.

The remarks under section 2 concerning the need for flexibility should be noted: standards are laid down in order to improve the productivity of the IT industry as a whole, not to stifle progress in its constituent parts. Thus, items that are identified as standards are a target of investment and will be widely available. They are mandatory in the sense that the providers and developers of tools must use them wherever it is sensible to do so. For the users of tools they represent the statement of the common research and development infrastructure that will be available to them and that should be used unless there are good reasons why not. These standards will be established by the Directorate through consultation with the R&D community.

In addition to these standards there will be many other research tools that are properly regarded as part of the infrastructure, which would benefit from being made more widely available, but to which it would not be appropriate to accord the same long term commitment as the standard infrastructure, e.g. compiler for a new experimental language. Such research tools may be developed rapidly or be superseded as new techniques are introduced.

It is essential that the infrastructure policy be evolutionary and not static. The first stage in successful exploitation of a new development will often be for it to move out of a

research project and receive the support necessary to become a component of the infrastructure

3.2 Systems Policy

Ideally it would not be necessary to have a policy for preferential use of certain computer systems. However the current state-of-the-art means that genuine portability of software is only rarely achieved, even at source code level, and therefore some degree of rationalisation is unavoidable if the best facilities are to be made widely available in a cost effective manner. Effective collaborative research is often dependent on the ability to rapidly bring new developments into widespread use within the R&D community. Without any policy on systems either the collaboration will not happen or research effort will be expended on tool replication. The market for Alvey funded systems will always be small and so the Directorate will always have to invest funds to support the replication of those new tools which it is not in the commercial interests of systems suppliers to support. Because the Alvey budget is small and finite the number of systems that can be supported in this way must necessarily be small. At a lower level, a policy on systems is necessary to achieve the benefits of a uniform supported service. The degree to which the availability of the essential functionality is constrained by hardware varies between Alvey subject categories and leads to some variations in policy. In each Alvey category however systems will fall into the following classes.

(1) Supported Systems

These systems are to be used on all projects where sensible, and will be eligible for Alvey funding. They will be the subject of direct investment by the Alvey Programme for mandatory and preferred standards and their use will be supported through the Alvey Infrastructure Programme (AIP).

(2) Specialist Systems

These are systems which meet some particular need not met by supported systems. They will be eligible for Alvey funding. They will not generally be the subject of Alvey investment in tools except on occasion in respect of their specialised facilities, nor will they be supported through the AIP.

The Alvey Directorate will where appropriate give encouragement for systems of UK origin to be brought up to supported status via the use of normal DTI funding mechanisms.

Where the use of other systems is permitted on Alvey projects they will not usually be eligible for funding. In assessing whether the computing facilities for a project are satisfactory the adherence to both the standards policy and the systems policy will be examined. The use of a system that jeopardises the maintenance of the standards policy will not usually be funded.

3.3 Policy Statements

Policy statements for each of the Alvey categories are contained in the appendices.

4. Infrastructure Programme

4.1 Background

Within the academic community it has long been recognised that considerable benefits accrue to research programmes from the provision of a supported national computing infrastructure. Such an infrastructure is the vehicle through which the infrastructure policy is executed. It brings a uniform level of service to a geographically dispersed community and has a powerful effect in enabling collaboration, improving productivity and concentrating investment in research tools. The SERC has run such services in all the component areas of computing infrastructure and the benefits have been amply demonstrated e.g.

—The Interactive Computing Facility of some 40 networked multi user minis for engineering research.

—Central mainframes running batch and interactive services, particularly for computationally intensive tasks such as high energy physics.

—The Common Base Programme supporting some 150 single user systems around the country.

As well as providing resources for established workers a supported service can be particularly useful in making 'pump priming' computing resources available for feasibility studies where the award of equipment is inappropriate.

The support provided for such a service falls into five areas:

(1) Systems support. By maintaining a single version of system software portability at the applications level is protected and standard communications facilities can be provided.

(2) Applications support. There is always a requirement for the continuing development of specialised tools and packages. Given a stable supported systems environment coordinated investment and distribution is possible.

(3) User support. This is an essential component in making the service accessible. Depending on the nature of the service this may be centralised (for mainframes) or distributed (for multi user minis).

(4) Operational Support. Where facilities are distributed experience has shown that a site manager, who can combine operational and user support, is highly beneficial in achieving the desired quality of service.

(5) Project Management. Such a service cannot be run without a significant level of overall management.

4.2 Alvey Infrastructure Programme

The Alvey Directorate will establish a support programme over all the components of the infrastructure for both industry and academia. A joint industry/academic infrastructure will be a new departure for the UK and should prove to be a major contribution to improving collaborative working and technology transfer. The degree of support provided to industry in the areas identified above will not be

the same as is appropriate in academia and may be expected to concentrate on the areas of application and user support.

The Alvey Directorate will set up an Infrastructure Steering Group comprising a representative of the Directorate and the community for each Alvey category. The Steering Group will report to the Board and will oversee the policy and its evolution. The Board will appoint a full time project officer to manage the day to day running of the Programme. The general support activities will be funded out of a budget contributed to by each category and administered by the Steering Group. Equipment for projects will be funded by the individual category budgets.

Where there is doubt as to whether a request for equipment on an Alvey project falls within the infrastructure policy then it will be referred to the Project Officer. The Steering Group will resolve disagreements.

Considerable work will be urgently undertaken by the Project Officer to formulate detailed plans for the whole programme. This section makes a number of general proposals, appendices contain initial plans for each Alvey category.

(1) Overall project management should come from the Alvey Directorate. The execution of specific projects should be contracted out to industry, government laboratories and universities/polytechnics.

(2) For single user systems support should be concentrated on sites by machine range. Coordination with the SERC Common Base Project is essential. Collaborative industry/SERC or industry/academic arrangements will be appropriate.

(3) For multi user minis there is already a project in place for IKBS and SE academic sites. The establishment of supported industry sites and any major expansion to cover other categories will probably be best served by regional support centres involving both industry and academia.

(4) For mainframes support will be localised at the service site, which is likely to be within SERC in the short term. User support for a much enlarged user community should involve industry.

(5) Communications. Alvey Net policy (JANET + PSS) is the essential basis of the service.

4.3 Software Development Contracts

Collaborative projects to develop and maintain application tools will be a normal part of the infrastructure programme and the conditions under which these projects should run must be considered. There are three reasons why these conditions may be slightly different to the general Alvey rules. Firstly, on development projects universities typically receive overheads which they do not get on research grants. Secondly, it may sometimes be the case that a particular development is required by the whole community but there is no obvious collaborator for the main producer of the tool. Thirdly, the requirements for dissemination and disclosure of infrastructure developments are not necessarily identical to those of other Alvey projects. The Project Officer will therefore establish general rules for these projects that will form the basis for agreements in individual cases.

5. Other Initiatives

5.1 SERC Common Base Programme

Within the academic community the SERC has established a Common Base Programme for all scientific and engineering subjects, not just information technology. The CBP provides a Common Base of hardware and software built around single user systems. This policy has been in operation for several years and is currently under review. A new specification has been drawn up and issued to a number of manufacturers as an operational requirement. A project has been set up to evaluate the responses to the OR with a view to bringing a new system in the CBP.

The Alvey Infrastructure Programme is not identified with the CBP but will obviously seek to maintain close links with it. In particular it has been decided that the CBP SUS evaluation project should be made the focus of an evaluation of systems for Alvey usage. The Directorate is coordinating the input of requirements and will be contributing resources to the evaluation project. It is recognised that the needs of the CBP and of the AIP will not necessarily be the same so a greater range of systems may be supported under the AIP than would be possible or appropriate under the CBP.

The term common base has passed into general use since the CBP was established; in this report the term is used only to refer to the SERC Common Base.

5.2 Advanced Workstation Architectures Project (AWSAP)

AWSAP is an Alvey collaborative project to design and develop a comprehensive, coherent architecture suited to the next generation of workstation or workcell which will be in use from the late 80s into the 90s in commerce and industry. The architecture may be proprietory to AWSAP but the objective is for maximum market adoption in the UK and overseas by workstation suppliers and users.

The project is to define common components of hardware and software, standards for user visible interfaces and internal interfaces, and the procedures and protocols for interconnection and interworking.

A three month design study is now in hand to prepare a technical specification and examine the technical feasibility of such a common architecture and its commercial justification. The study will also determine how the main project should run and it is expected that many more companies, possibly over 20, will take part in a two to three year programme.

This project is taking a 'clean sheet' approach and may be regarded as a category C user of the infrastructure. Its progress will be closely monitored by the Infrastructure Steering Committee.

5.3 Advanced Communications

The Infrastructure Director intends to support certain research projects concerned with advanced communications. These projects will be set up on a pilot high speed network established in parallel with the service network. It is not intended that the Alvey R&D programme should go forward on the pilot service.

Appendix A

Communications Policy

1. Introduction

The communications policy is concerned with three main areas:

(1) Open systems interconnection, local and wide area.

(2) Distributed systems, local and wide area.

(3) Personal communications and document exchange.

The standards and infrastructure programme are presented separately for each of these areas in sections 2-4. Section 5 concerns links with ESPRIT. The general communications policy of the infrastructure Director is included in section G.

2. Open Systems Interconnection

2.1 Wide Area Networks

2.2.1 Standards

For supported systems the mandatory standards are as follows:

(1) X25 levels 1,2,3, PSS compatible.

(2) Yellow Book Transport service as defined for X25 in Annex 1 of [1].

(3) X29 and TS29 (Green Book) Remote Terminal Login. Single user systems will not be expected to support incoming calls. Ref [2].

(4) FTP80 (Blue Book) File Transfer protocol. Ref [3].

(5) Mail (Grey Book). Ref [4].

These protocols are all currently in widespread use on the Joint Academic Network (JANET), are available for the IKBS and SE supported systems and are therefore the only available basis for the immediate support of services.

2.1.2 Support Programme

(1) All supported systems will be provided with the standard communications.

(2) All infrastructure sites (i.e. those with a supported system on an Alvey project) will be networked. The cost of this will be met from the Steering Group budget contributed to by each Alvey category.

(3) The network service will be supplied by existing networks, by enhancement if necessary. PSS and the JANET will be constituent parts of Alvey Net, which is a management entity rather than a separate network. An assessment of expected network load and capacity is an urgent priority of the Alvey Net management. The cost of network enhancement will be met from the Alvey Infrastructure budget.

(4) The Alvey Directorate wishes to actively promote the migration of its network services to ISO/CCITT standards and will cooperate to this end with ITSU, the DTI and the management of the network providers in establishing an action plan to achieve this. However, a complete set of ISO Standards will not be fully defined before 86/87 and a piecemeal transition of protocols is judged not to be a viable option.

2.2 Local Area Networks

2.2.1 Standards

Reflecting existing investment in the academic community in LANs two technologies have standard status for supported systems:

(i) Cambridge Ring to CR82 specification. Ref [5].

(ii) CSMA/CD providing the IEEE 802.3 MAC service.

For Cambridge Ring the standard protocols up to transport service are defined in [6].

There are two main options for LAN OSI protocols for CSMA/CD:

(i) LLC2 subset of IEEE802.2 plus X25 level 3.

(ii) LLC1 subset of IEEE802.2 plus ECMA 72 Class 4 transport service.

Both these options are currently part of the ITSU intercept strategy. The Joint Network Team of the Computer Board and Research Councils is studying the issues and hope to make a policy decision for its own programme in 3Q84. It is expected that this policy will become a mandatory requirement for Alvey supported systems, mandatory for academic infrastructure sites (unless served by X25), and a preferred standard for industry infrastructure sites. It is not yet clear whether it is technically feasible to run the two different options with the standard WAN protocols between.

2.2.2 Support Programme

(1) All supported systems will be provided with standard communications in at least one technology.

(2) Where more than one machine is provided to an infrastructure site LAN facilities will in general be provided. This may be via X25 if that is a cost effective solution or through the standard LANs. Funding will be as for WAN connections, and will generally be restricted to connection costs rather than complete LAN installation.

(3) As for WANs the Directorate will promote migration to ISO/CCITT standards.

3. Distributed Systems

3.1 Standards

OSI protocols apply between what are known as 'end systems'. An end system may be composed of systems linked together in any way and such linking is the subject of distributed systems. The protocols that are used to link systems across a network into a single distributed system are currently an active area of research. OSI protocols represent a conservative approach and are not intended to tackle distributed system problems, such as file servers, remote paging, etc. They do not allow various advantages of LANs to be exploited e.g. speed, reliability. More LAN specific protocols could be employed for high speed intermachine interaction (e.g. remote process execution). There are no protocols that can currently be given standards status.

In certain areas of the programme (IKBS and MMI in particular) a strong requirement for interprocess communication facilities has been identified as part of the UNIX environment. As an interim position the Berkeley 4.2 IPC will be taken as a preferred standard definition of IPC facilities. A mapping of this IPC onto the standard LAN protocols is urgently required. It is recognised that migration from this interim position is inevitable.

3.2 Support Programme

(1) Supported SUSs and MUMs for IKBS and MMI will be provided with a UNIX system supporting Berkeley 4.2 style IPC.

(2) Activities to define interim remote procedure call (RPC) standards will be supported within the limits of available manpower. In the longer term the definition of ISO standards is desirable.

4. Personal Communications and Document Exchange

A three tier strategy will be followed comprising:

(i) personal message passing via electronic mail and conferencing

(ii) document preparation and exchange by word processing plus Teletex

(iii) document preparation, exchange and high quality output for UNIX systems.

4.1 Personal Message Passing

(1) Grey Book Mail is the mandatory standard protocol between mail servers. A supported service is being set up on two mail servers located at NPL and RAL. Mailboxes will be free to Alvey participants but terminals and connection are at the participants expense or funded on projects. Full access details will be made available shortly.

(2) A number of conferencing and bulletin board facilities are under trial use by the Alvey programme and there is a commitment in principle to making some such service available.

4.2 Document Exchange—non UNIX

The electronic mail service is not designed for document exchange or for the support of word processing. Many Alvey projects will have access to UNIX facilities and will therefore make use of the services described below in 4.3. A gap exists, however, where a networked UNIX system is not available and it is proposed to fill this by support for a word processing plus Teletex (S.62) facility. The support will take the form of Alvey action to make suitable products available on the market.

4.3 Document Exchange and High Quality Output

Where UNIX is available AT&T System V version 2.0 Documentor's Workbench troff (based on ti-troff) will be the preferred standard for document preparation and Mail or FTP80 will be used for document exchange. As part of the proposed mainframe UNIX service it is hoped to be able to handle troff files from any Alvey project.

5. ESPRIT

The Directorate will seek in all areas of its communications policy and programme to achieve interworking and convergence towards common standards with ESPRIT. Initial progress will be in the area of personal communications through mail convertors and common conferencing facilities. Full details of the programme have yet to be worked out.

6. Communications Policy

Objectives:

—To provide common electronic communication services to the Directorate, the government and academic agencies involved in the Alvey programme, all commercial participants who wish to take advantage of the service and approved additional users having some relevance to the Programme.

—To promote development of advanced practical systems and stimulate research to this end.

—To support state of the art activity by using and encouraging use of advanced techniques and components wherever feasible.

—To foster the use of emerging standards where these have been agreed and are adequate to support practical systems, and to promote continuing development of network and communication standards generally.

—To stimulate and sponsor research in support of these objectives in particular and other communication areas meeting the Alvey criteria in general.

Policies

1. To keep constantly under review through the Alvey Directorate the requirements of all the Alvey participants for communication services, develop systems to meet common requirements and publish a programme for their provision.

2. To use PSS to provide basic connectivity.

3. To use PSS to support an electronic mail service. This will be based initially on 50 industrial user mailboxes on an existing server at RAL, but a new similar system at NPL will be ready by June 84. This will support some 1500 users. These two mailservers, and a third at RSRE will intercommunicate using the JNT Greybook protocol.

4. To use the Blue Book File Transfer protocol as Alvey file transfer standard.

5. To actively promote adoption of ISO and CCITT standards, in particular by cooperation with ITSU of DTI, and generally by use where feasible in services and sponsored projects.

6. To encourage proposals for development of practical standards conversion capability at any level.

7. To complete a study to propose a design and estimate the costs of an Alvey Communications Environment comprising a wide bandwidth network of switched megabit links between selected centres and participants sites.

8. To establish a network on the basis of the study report.

9. To sponsor network and communication research projects, particularly in association with the network to be provided.

10. To accept some research risks in the network to be provided in the interests of providing services in parallel with continuing research.

7. References

1. Transport Service: Yellow Book
2. TS29: Green Book
3. FTP80: Blue Book
4. Mail: Grey Book
5. CR82 Interface Specifications (Orange Book)
6. CR82 Protocol Specifications (Orange Book)

Appendix B

Software Engineering Infrastructure Policy

1. Introduction

This paper is a modified version of the current Alvey SE Infrastructure Policy as given in "Alvey Programme, Software Engineering Strategy, Nov 83".

2. Proposed SE Infrastructure: Summary

2.1 Software

SERC Common Software Base (see appendix 2) to be used for:

(a) all academic projects (where sensible)

(b) all 1st generation IPSE consortia

(c) all other SE consortia where sensible.

2.2 Single User Systems

SERC Common Hardware Base (see Annexe 1) to be used for:

all academic projects (where sensible).

Alvey consortia to conform to Infrastructure W P recommendations (where sensible).

2.3 Multi-User Minis

SE Infrastructure Policy (see Annexe 1) to be used for:

(a) all academic projects (where sensible)

(b) all industrial projects (where sensible)

2.4 Mainframe

New mainframe service to be set up based on IBM compatible CPU running UNIX and Common Software Base.

2.5 Communications

All Alvey SE participants to be networked together immediately by PSS or JANET to run

(a) coloured book electronic mail (local or via RLGB)

(b) COM

(c) coloured book file transfer.

Long term Alvey communications to evolve to use ISO standard protocols.

1. Introduction

The Alvey Software Engineering programme is establishing an infrastructure, containing the SERC's Common Base Policy (CBP), which will form the framework within which to research and develop the first and second generation IPSEs (integrated project support environments) as outlined in the Alvey Software Engineering Strategy.

In outline the envisaged SE Infrastructure will see a typical project have a powerful, multi user mini (MUM) Unix system linked by high speed local area network to a set of high performance Single User UNIX Systems (SUS) running the Common Software Base and Communications protocols with access to the large data storage and cpu power of a mainframe also running UNIX and the Common Software Base.

2. Summary

2.1 Software

a. SE Research Tools (examples)	ML, HOPE, LCF. Ada, Modula-2. Affirm, Stanford Pascal Verifier. Boyer-Moore, Iota.	
	LISP	(IKBS recommended version)
	Prolog	(IKBS recommended version)
	Pop-2	(IKBS recommended version)
b. Languages	Pascal	(ISO Standard)
	Fortran 77	(ANSI Standard)
	Cobol	(ISO Standard)
	Ada	(ANSI Standard)
c. Graphics	GKS	(BSI and draft ISO Standard)
d. Operating System	UNIX	(32 bit virtual memory)

2.2 Computer Systems

a. Single User Systems	32 bit SUS	(current CBP m/c is PERQ)
b. Multi-user Mini	32 bit MUM	(VAX, GEC Series 63)
c. Mainframe		(to be decided)
d. Servers		(Not yet available)

2.3 Communications

a. Local Area Network	Cambridge Ring	(UK CR82 standard)
	Ethernet	(CSMA/CD for IEEE802)
b. Wide Area Network	X25	(JANET & PSS)
c. Protocols	JNT Coloured Books	(UK academic standards)

A more detailed exposition of the technical components and philosophy of the policy is given below.

3. SE Research Tools

The following list contains those tools which are known to be of interest to several groups of researchers and therefore would benefit from wider availability via the SE Infrastructure. The list should be regarded as 'provisional'.

3.1 Functional Languages

ML, HOPE, LISP.

An STI workshop at RAL led to the expression of considerable interest by the community in having LISP, ML and HOPE running identical versions on all SE Infrastructure supported systems.

ML is associated with LCF (see below).

3.2 Logic Languages

PROLOG

There was some interest expressed at the RAL workshop in PROLOG. The SE programme should be able to get PROLOG via the IKBS programme.

3.3 Procedural Languages

ADA, MODULA-2

Ada is available on VAX/UNIX from the SERC funded work at York. The same compiler is being moved to the PERQ.

The Cambridge VAX/UNIX Modula-2 compiler is being moved to the PERQ.

3.4 Verifiers, Theorem Provers

LCF, AFFIRM, SPV, B-M, IOTA

The UK does not have all of these systems. Only LCF and Boyer-Moore are available, currently on the Edinburgh DEC-10 which will be shutdown in October 1984. LCF is also available on VAX/UNIX.

There is considerable support for having LCF available on all SE Infrastructure supported systems.

The Stanford Pascal Verifier is being mounted on the VAX/UNIX by Strathclyde via an SERC EMR.

The Alvey SE Director is working at obtaining AFFIRM and IOTA.

4. Languages

4.1 Cobol, Fortran 77, Pascal, Ada

Cobol, Ada, Pascal and Fortran 77 have been chosen as they are the most popular languages. They possess the properties of portability and official standard definitions. There is a large amount of software already written in them which allows people to make use of existing investment.

There is considerable SERC support for Fortran 77 and Pascal. This takes the form of software tools and techniques developed by the SERC Software Technology Initiative and the activities of the SERC Computing Service team and the SERC Common Base Programme.

The technical definition of Pascal is given in (ref 1).

The technical definition of Fortran 77 is given in (ref 2)

The technical definition of Cobol is given in (ref 16).

The technical definition of Ada is given in (ref 17).

4.2 Other Languages

Other languages will be available with the set of software tools in the SE Infrastructure. For instance the Unix 'C' language is already available. LISP and Prolog are being implemented.

These other languages will not receive the same degree of support and tool development as Pascal and Fortran. They are not 'blessed'. This situation must be reviewed regularly.

Evolution of status from 'other' to 'blessed' is possible.

4.3 Mixed Language Working

It is a requirement of the SE Infrastructure that 'blessed' languages should be inter-workable at the procedure call level i.e. a Pascal program can call a Fortran subroutine which can call a Pascal procedure, etc. This is a vital capability to ensure maximum use of standard components. It is ridiculous to have to, say, reimplement a Fortran graphics package in Pascal because Pascal cannot call Fortran.

Interworking has implications for compiler construction and operating system development. It has its limitations and difficulties, e.g. the difficulties in enforcing type checking across procedure interfaces, but its benefits outweigh its drawbacks. (Reference 15.)

5. Graphics

In line with the policy of supporting international standards and portability aids the SE Infrastructure has 'blessed' GKS 7.2 (the ISO Standard Graphical Kernel System) as its basic graphics package. GKS will be available on all SE Infrastructure supported systems to help the transfer of graphics software and, via metafile standards, pictures themselves.

The technical definition of GKS is given in (ref 3).

The technical definition of metafile standards is under development by ISO.

6. Operating System

6.1 CBP UNIX

Unix is already a de facto standard in many academic institutions in both USA and UK. It has enabled a great deal of software to be shared amongst research groups and has built up a large quantity of widely applicable software.

Unix is being used increasingly by industry again both in the USA and UK. The CBP philosophy is based on the following properties of Unix.

a. It is popular, i.e. a de facto standard.

b. It is implemented on a wide variety of makes and sizes of computer (IBM 370-M 68000).

c. It is cheap ($150 per PERQ).

d. It has a large body of user level software.

e. It is used by both industry and academia.

The CBP Unix has the following properties

a. It is full '32 bit'.

   Arithmetic is 32 bits by default to overcome the annoying limitations to microprocessors. 8,16,32,64 bit quantities are available.

b. It is virtual memory.

   Full 32 bit addressed linear address space (via paging) removes the size restriction which is often so frustrating.

The technical specification of Unix is given in (ref 4).

6.2 UNIX Evolution and Standards

There are several versions of UNIX either in existence or soon to be announced. These include Berkeley 4.1 and 4.2, Bell version 7, System III and System V version 1.0 and 2.0.

The SE philosophy is to run the same, stable version of UNIX on all the different types of hardware supported by SE i.e. only one version of UNIX will be supported by SE.

There must be a balance between the benefits of new developments and the benefits of stability and standardisation. Thus moving to a new version of UNIX will be a major evolutionary step for the SE Infrastructure, especially if and when more than one CPU type is involved.

For supported systems UNIX is mandatory. Ideally the standard should be the AT&T version. However, the support of languages and environments and the interactive style of program development make particular demands that in the past have been met only by the Berkeley version:

—virtual memory supporting incremental compilation, non-local patterns of access, etc,

—interprocess communication,

—superior job control and user environment.

AT&T System V2.0 is a significant move towards these requirements but is known not to provide IPC, and the initial version of virtual memory is not yet known to be satisfactory.

It is unavoidable therefore that a certain degree of deviation from standard AT&T UNIX will be required at least on the SUS and MUMs. This is likely to limit the number of different systems that can be brought into the list of supported equipment since resources will be required to enhance the basic systems.

The policy will be to converge on the AT&T version whenever that is possible.

7. Computers

7.1 Single User Systems

The PERQ was the first machine which satisfied the requirement for a high performance single user system (see Appendix 2). Other machines are likely to follow (some are already here). The expected proliferation of machines will tend to fragment the software development activities because some things will always be machine specific. The Directorate therefore wishes to balance the benefits of standardisation (which acts against change) with the need to give state of the art facilities to researchers (which requires change). The SE Infrastructure is therefore expected to include more than just today's PERQ but such changes must be taken infrequently and given very careful consideration beforehand.

It should be borne in mind that the dominant criterion for choosing a single user system must be that it runs the Common Software Base rather than has some new hardware feature. The investment in software is already so large that computers must be purchased which run existing software rather than the Directorate's money be wasted on reimplementing existing software on some new hardware. Manufacturers will have to understand the changing balance of power between them and their customers. The cpu independence of Unix is a key factor in this equation.

The recommended CBP PERQ configuration is:

PERQ: 2 Mbyte main memory
16K writeable control store
35 Mbyte disk
tablet
puck (3 button 'mouse')
LAN interface
X25 front end (one per installation for connection with WAN)
1 Mbyte floppy disk
100 Pixel per inch A4 display

For advice on peripherals such as printers suitable for use with PERQ contact CBP User Support at RAL.

The technical specification of the PERQ is given in (ref 5, ref 6).

7.2 Multi User Mini

The IPSE requires not only high performance SUS for highly interactive tools but also major computational power (verifiers), storage (database, backup) and peripherals (printer, communications, tape deck, archiving). It is expected that each research project engaged in a significant amount of SERC SE research will have a MUM to complement the SUSs. The MUM must run the same Common Software Base as the SUS. The MUM is a necessary infrastructure component because

a. SE investigators need to run imported software

b. a MUM allows more users to be given access to simple services which do not require a PERQ

c. SUSs need a host machine to provide an interim 'server' capability for archiving, printing, etc.

d. many SE investigators now require computationally intensive tools such as theorem provers, especially the VAX/UNIX implementation of LCF from Cambridge

e. it is important, as far as is possible, to provide technical compatibility with the IKBS and other Alvey communities as productive cross fertilisation is likely

f. an infrastructure machine could provide a pump priming facility to enable work to be done prior to and in preparation for an Alvey grant.

A typical SE Infrastructure MUM configuration is

MUM: Powerful, multi-user minicomputer e.g. GEC series 63, VAX
4-8 Mbytes main memory
500-1000 Mbytes disk storage
9 track 1600 bpi tape deck
Line printer, upper and lower case
Documentation quality printer (till Laser Printer LAN server)
Local area network connection
Wide Area network connection (till LAN server)
5 'conventional' VDUs (+5 SUSs)

7.3 Linked SUS-MUM IPSE Distributed Computing System

The MUM/UNIX system is envisaged as being accessible to 5 users via conventional terminals and to 5 users via high performance SUS PERQ/UNIX systems, via a high speed local area network. Peripherals, including a laser printer and wide area communications, will be shared by all (SUS + MUM) systems through LAN servers eventually and via the MUM in the short term.

Such a distributed computing system will consist of 6 UNIX systems connected via file transfer and remote login facilities based on the JNT coloured book protocols. It is envisaged that a more unified, logically integrated (remember the IPSE target), physically distributed UNIX software environment will come into operation in the mid 80s as a result of non UK work, the SERC's DCS Programme and general UK progress towards the IPSE.

7.4 Mainframe

SE projects are likely to need

(a) large batch CPU power
(b) fast response + CPU power
(c) large data storage
(d) general service
(e) pump priming

Large batch CPU power is likely to be required for work involving mechanical theorem proving. This type of work tends to be computationally intensive. For instance, ML/LCF is sometimes run overnight on a VAX 750 to reach a conclusion.

Similarly some theorem provers are interactive, using the man's skill to help speed up the proof process. Such systems still require large CPU power, greater than a minicomputer can provide.

Large data storage is likely to be required for archiving MUM and SUS data, for holding Alvey-wide data such as software libraries, conferences (in the electronic mail sense), online news bulletins, lists of projects and other management information.

Large data storage combined with CPU power is likely for tools which require both such as analysers (control + dataflow) and quality certification.

There is a requirement for a general service to researchers whose departmental facilities are inadequate. Often it does not make sense to buy a small VAX for a small project if a mainframe service exists. Pump priming is an extremely useful facility (cf ICF DEC-10) but cannot reasonably be done on anything but a mainframe.

A mainframe service is also a good coordination vehicle because software and data can be shared between many people. In particular the DEC-10 at Edinburgh has shown this to be true; additionally the DEC-10 has had the advantage of facilitating the development of single versions of software tools, whereas multiple MUMs and SUSs act against this uniformity. The mainframe version is a natural reference standard.

Much SE work is highly interactive in nature. A mainframe service for the most demanding users will therefore only be of benefit if the networking service is of a high bandwidth (greater than 4.8 Kbit/second) and the system response time over the network is low (less than 0.2 seconds). Providing such a service may therefore entail investment in high performance network links. Ideally the development of this service will include provision of the 5620 type of bit map displays, but current network protocols are particularly weak in the area of supporting highly interactive work.

A mainframe service only makes sense if it runs the same basic software as the SUSs and MUMs. Thus the SE Directorate would like to have a large mainframe, such as an ICL Atlas 10 or an IBM 3083, running the Common Software Base. This is feasible today through Amdahl's UNIX system UTS which runs on both the Atlas 10 and the IBM range.

8. Local Area Network

8.1 Cambridge Ring

The SE Infrastructure requires a fast local area network to link its machines together. The Cambridge Ring has been chosen because it is

a. a UK draft BSI standard (CR82, ref 7).

b. SERC CBP and SERC DCS Programme' common equipment

c. has protocols already implemented for Unix which are a UK standard

d. it is an easily purchased and maintained commodity from a variety of UK suppliers.

The technical specification of the Cambridge Ring is given in (ref 7) and the protocols in (ref 14).

8.2 LAN Evolution

The Cambridge Ring is not the only LAN currently available, but has been chosen as the CBP LAN for the above reasons.

There are several different types of Ethernet and Token Ring LANs available or soon to be announced. The IEEE 802 standard initiative is having a beneficial influence but has yet to be adopted as an ISO standard.

The SE Infrastructure will support CSMA/CD to IEEE 802.

8.3 Campus X25 Switches

Where a campus has installed an X25 system to act as a LAN then the SUS can access this via the hardware and software given under section 8, i.e. X25 campus LANs are 'blessed' by the CBP and the SE Infrastructure.

9. Servers

The long term objective of the SE Infrastructure is to exploit the advantages of distributed computing and LANs which can be realised as Servers. The following Server requirements can be identified as desirable but not yet deliverable as 'service' equipment. There is an urgent need to develop such servers into commercial products.

9.1 Publication Quality Printing

There is a requirement for sophisticated, high quality (at least 300 pixels per inch) text and graphics printing capability to complement the Single User System's display.

Examples are hardcopy of scientific papers (camera ready including diagrams), graphical software tool output, 'mathematical' text (proofs) and so on.

It is envisaged that this need will be met by small, relatively cheap (£5K) laser printers, one per department, configured as a LAN server. Until this technology is readily available (1984?) such items as Diablo daisy wheel printers and Versatek graphics devices are suggested (Contact RAL CBP team for advice).

9.2 LAN/X25 Gateway

It is seen that an LAN to X25 (SERCNet and PSS) gateway will be the most cost effective way of connecting a number of machines to the WAN. No products are currently supported.

9.3 LAN/LAN Gateway

Multiple campus LANs, whether the same type or not, are likely to arise with the consequential need to connect one LAN to another.

No products are currently supported.

9.4 File Server

Single user systems cannot hold all of the data to which a single user requires access. Nor can a SUS handle file backup and archiving requirements.

In the short term the SE Infrastructure recommends that SUS are not used stand-alone but are connected to multi-user machines with suitable peripherals to allow file access and archiving.

The more desirable solution is to have file and/or archive servers. No products are currently supported.

10. Wide Area Network

The SE Infrastructure requires a national wide area network to link both people and machines. The network will follow Alvey Communications Policy. The current SE Infrastructure uses JANET and PSS which are technically compatible X25 networks linked by a gateway.

The SE Infrastructure also requires access to Europe (including Scandinavia) and the USA. Such links are not all easily available.

The PERQ-X25 and VAX X-25 connections, in the short term, will be via the York LSI-11 transport service front end originally designed for the PDP-11.

The technical specification of JANET X25 network is given in (ref 8).

11. Protocol Strategy

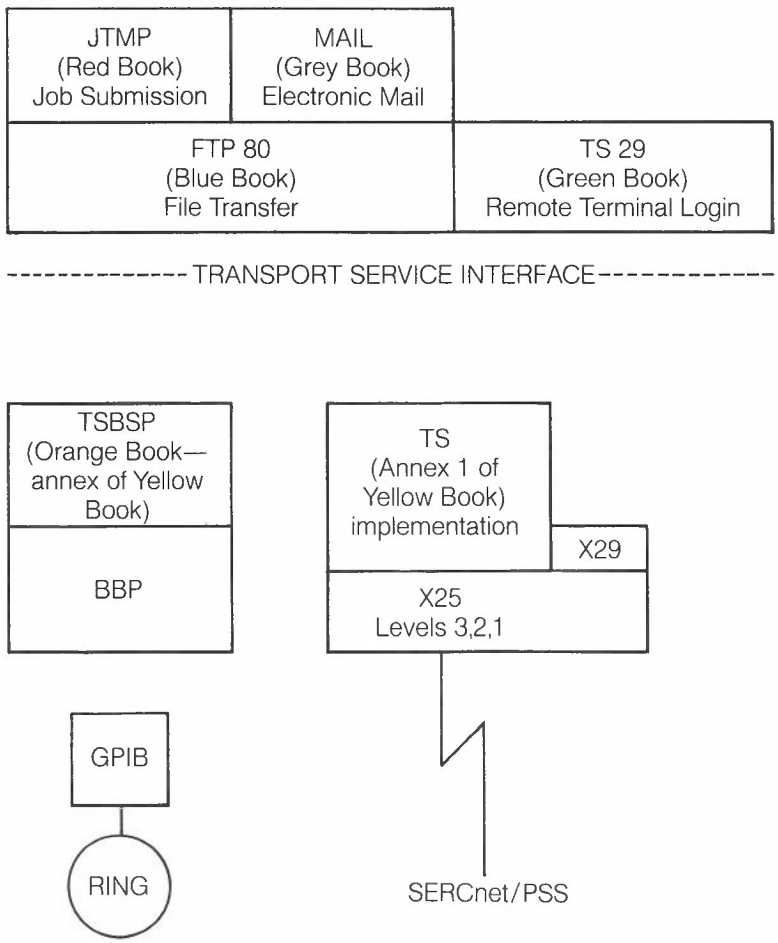

11.1 CBP Protocols

The protocol strategy is based on the de facto UK academic standards approved by the SERC/CB JNT in their 'coloured books'. The adoption of the Wide Area Network protocols of transport service and above for the local area network use gives a useful unification of LAN/WAN facilities. The average user sees only one and the same mechanism to move files, mail, etc. between machines independent of distance (i.e. local or wide area net). The adoption of transport service also gives a degree of hardware independence for local area networks.

The use of wide area protocols for local area networks is 'conservative' in that it does not allow various advantages of LANs to be exploited e.g. speed, reliability. More LAN specific (light weight) protocols could be employed for high speed intermachine interaction (e.g. remote process execution). Such protocols should only be 'blessed' if they attain a measure of widespread acceptability. Specific research projects are likely to require lightweight protocols. They should not be discouraged in appropriate circumstances.

Transport Service around the Ring is implemented by TSBSP (Transport Service Byte Stream Protocol) running above BBP (Basic Block Protocol). These are the de facto UK academic Ring standard protocols based on Cambridge University's work.

The protocols specifications are given in (refs 9-15).

11.2 Conferencing, Bulletins

Electronic Mail as implemented over the Grey Book is an extremely useful facility. However, experimental work at various sites in the world has shown the potential advantages of more sophisticated facilities above simple mail. Such facilities include message based conferences and public electronic bulletin boards.

No ISO approved or de facto standards exist in these developing areas. The SE Infrastructure could possibly evolve to include such facilities. Esprit is using COM.

11.3 Protocol Evolution

The JNT coloured books and the CR82 Ring protocols are not ISO standards nor are they likely to be. It will be necessary eventually to change the protocols on both WAN and LANs in the light of current development work on protocols to whatever emerge as international standards. This will be a major change for the entire network community and will not come quickly.

12. Portability

Fortran 77 and Pascal will, for example, allow PERQ CBP software to be moved to and from other non PERQ computers. However it is recognised that even when programs are written in Fortran 77 and Pascal much work often has to be done to move them because of the inbuilt operating system dependencies. By using 32 bit, virtual memory Unix as a de facto standard execution environment it should be much easier to move programs in Pascal and Fortran 77 from one CBP Unix system to another.

Portability is also one of the reasons for backing national and international standards generally, hence the use of the GKS graphics package. GKS will be available on all SERC supported machines.

Portability of software is also one of the aims of the networking side of the SE Infrastructure. Good communications are needed if software is to be easily shared by geographically dispersed research groups.

13. Applications Specific Support

The SE Infrastructure is expected to be expanded to include some items related to specific applications. These might possibly be the NAG library, RAL graphics library, etc. as well as software development tools from the SERC CBP, Alvey SE, IKBS, etc. In addition much applications specific software will be generated 'on top of' the SE Infrastructure and which will be generally available but which will not actually be part of the SE Infrastructure. The SE Infrastructure is supposed to form the 'base' not the totality of available software.

14. General Points

a. Great stress should be laid on the fact that the SE Infrastructure does not see single user systems as standalone systems. Networking is the key to file backup, mail, software update and interchange.

b. SE Infrastructure links people just as much as computers.

c. SE Infrastructure aims to back international standards if possible.

d. Software sharing and portability only really comes when both the programming language and execution environment (i.e. operating system) are defined. The corollary is "it's OK to change the machine—just don't change the (user/program and program/operating system) interfaces".

15. References

1. ISO Standard Pascal (BS 6192)
2. Ansi Standard Fortran 77
3. GKS draft ISO standard
4. UNIX Manual
5. PERQ glossy
6. PERQ hardware manual
7. CR82 UK Ring hardware specification
   CR82 Interface Specifications
   Orange Book

8. SERCnet X25 specification

9. TS29: Green Book

10. FTP80: Blue Book

11. JTMP: Red Book

12. MAIL: Grey Book

13. Transport Service: Yellow Book

14. TSBSP, BBP
    CR82 Protocol Specifications: Orange Book

15. Mixed Language Working
    A Williams
    RAL

16. Cobol Standard

17. Ada Reference Manual

.....

Annexe 1 (to Appendix B)

SERC Common Base Policy

1. Introduction

1.1 Overview of Distributed Interactive Computing

The appearance in the market place of cheap high powered single user computer systems with good interactive capabilities via high precision displays, linked together by high speed local area networks, heralds a completely new way for most SERC Investigators to achieve the major part of their computing requirements.

Within the next few years, many such systems will be available from different manufacturers. Consequently there is a likelihood of many different systems being purchased in the SERC environment leading to a great deal of duplication of basic software development.

SERC sees a need for a coordinated development plan to ensure that the UK makes the best use of its finances and of its limited manpower. The SERC has therefore decided on a strategy of creating a common hardware and software base for software development which will encompass all scientific subject areas. Briefly the common software base will be Pascal and Fortran running under the Unix operating system implemented on the common hardware base of PERQ single user computers linked locally by Cambridge Rings and nationally by the X25 wide area network systems (SERCnet and PSS).

SERC Subject Committees will participate in the implementation of this policy by enabling central purchasing of PERQs for grant holders to be done via Central Computing Committee and by ensuring that investigators use the PERQ in all appropriate circumstances as well as encouraging them to follow the common base software development policy. The Common Base Policy is not the same as standardisation, however, and it will evolve as the state of the art improves.

1.2 Common Base Policy

The whole academic community, not just Computer Science, is a major user and developer of software and so the degree of ease with which software can be developed affects the scientific productivity of many researchers.

The SERC has approved a plan to increase the productivity of scientific research requiring computing by:

(1) facilitating scientific cooperation by:

  (a) person to person links
  (b) computer to computer links
  (c) common software and hardware base policy.

(2) Set in motion a coherent plan to exploit software tool production by making such tools/techniques widely known and available in forms which can be readily used by the whole user community.

Currently the academic software technology base is very non-uniform in that the knowledge, experience, tools, techniques and equipment vary considerably between projects. The motivation to create a common Hardware and Software Base is to bring together all of the best existing tools, packages and techniques into a uniform framework so that the 'whole' is more effective than the 'sum of diverse parts'. This will be achieved via EMR contracts to move existing software into the common base, specific purchases, the direct results of SERC research projects using the common base equipment and the 'snowball' effort that will be generated as a natural consequence of providing a state of the art hardware base. A good example of the common base 'snowball' effect is the widespread use of the Unix operating system which has enabled a large number of software tools to be made available throughout the UK academic community.

The Common Base Policy briefly is:

(a) common software base,

(b) common hardware base,

(c) common communications.

The SERC wish the common software base to be the Unix operating system and the common hardware base to be the PERQ. The PERQs should be networked together via Cambridge Ring, SERCnet and PSS to allow widespread cooperation between users and developers. This combination of software and hardware is widely accepted as being the best combination for developing software in the coming years. A common base does not imply rigid standardisation however.

Computer technology develops at a rapid pace and it is expected that the next few years will see the cost of single user systems decline and their quality and capability increase. Therefore today's PERQ is seen as only the first machine forming the common hardware base. The common base will develop over the coming years.

1.3 Outline

In outline the Common Base Policy comprises

a. Pascal (ISO Standard)

b. Fortran 77 (Ansi Standard)

c. GKS (BSI and ISO Standard)

d. UNIX (32 bit virtual memory)

e. PERQ (High performance single user system)

f. Cambridge Ring (Local Area Network)

g. X25 (Wide Area Network)

.....

Annexe 2 (to Appendix B)

The PERQ

The PERQ is a high powered, single user computer system with a high precision display system which provides a significant improvement in the quality and speed of interaction. Its main features are:

(1) *High Speed Processor*
Approximately 1 million 'high level' machine instructions per second giving around two-thirds the CPU power of a VAX 11/780. The CPU is micro-programmable for further speed gains.

(2) *High Quality Display*
A4 size, 1024 × 768 pixel, high resolution black and white display featuring 60Hz non-interlaced refresh rate which enables pictures to be moved cleanly and rapidly as well as giving a significant improvement in the clarity of text and diagrams equal to a printed A4 page.

(3) *User Friendly I/O Devices*
A 2-D tablet and voice synthesiser, allied to the high quality screen, enable a much improved man-machine interface to be created.

(4) *Large Virtual Memory*
A 32 bit address virtual memory system.

(5) *Local Filestore*
A 24 Mbyte Winchester disk and 1 Mbyte floppy give a single user a large amount of local storage capacity.

(6) *Fast Communications*
Local communication at 10 Mbits/sec via Cambridge Ring and Ethernet. Standard RS232 serial and IEEE 488 parallel interfaces are also provided.

A high quality, superbly interactive computing system is created if each investigator has his own single user PERQ linked to his colleagues' PERQs and other departmental computing resources by a Cambridge Ring, with inter-university cooperation being fostered by the National X25 network connections.

Appendix C

IKBS Infrastructure Policy

1. Standards

1.1 Operating System

For supported systems UNIX is mandatory. Ideally the standard should be the AT&T versions. However, the support of IKBS languages and environments and the interactive style of IKBS programme development make particular demands that in the past have been met only by the Berkeley version:

—virtual memory supporting incremental compilation, non-local patterns of access, etc,

—interprocess communication,

—superior job control and user environment.

System V2.0 is a significant move towards these requirements but is known not to provide IPC or virtual memory to the functionality required.

It is unavoidable therefore that a certain degree of deviation from standard AT&T UNIX will be required at least on the SUS and MUMs. This is likely to limit the number of different systems that can be brought into the list of supported equipment since resources will be required to enhance the basic systems.

The policy will be to converge on the AT&T version whenever that is possible.

1.2 Languages

1.2.1 PROLOG

The Edinburgh DEC10 system is a de facto standard already endorsed by the community. It will be taken as defining the mandatory standard for all projects producing PROLOG tools and environments.

Standards activity in the BSI/ISO sense has been started even though it is recognised that further development of PROLOG will occur.

A study has been commissioned on the future directions that Prolog developments should take to enhance its acceptability as a tool for commercial use.

1.2.2 LISP

In framing a policy on LISP we must take into account that several of the manufacturers who are likely to be major suppliers of systems running Lisp already have clear commitments to the different families of Lisp e.g. Xerox-Interlisp, DEC—Common Lisp, Hewlett-Packard—Portable Standard Lisp. We have to recognise that we shall not have any short term influence over these commitments and we must therefore address the issue of portability between dialects. A methodology is being developed (known as Minimal Lisp) which will greatly assist the porting of Lisp code between dialects. This methodology will be defined by the end of 1984 and will be a mandatory standard for all projects producing tools in Lisp.

The entrenched positions of manufacturers means that the choice of supported machines is likely to have a significant influence over any choice of a standard Lisp implementation. However, if there is to be such a standard, entailing as it does a commitment to make it available on all supported systems, there must be no technical or licencing obstacles to its widespread use. It must not be too expensive in manpower terms to port it, nor must the UK find itself at the mercy of US export and licencing restrictions.

The cost of porting Interlisp means that it cannot be taken as a standard and there will be no commitment to making it generally available to the IKBS programme.

Common Lisp is likely to be widely used and the ability to readily import code written in it is likely to be very advantageous. Franzlisp, which belongs to the same general family, is currently the preferred standard Lisp.

The UK is in the fortunate position of having its own Lisp expertise built up around Cambridge University Lisp (CULISP) which is very close to PSL. It is potentially possible therefore that by means of compatability packages on top of CULISP it will be possible to support at least some clean subset of code from Common Lisp and Interlisp.

The Directorate will actively review the technical and licencing aspects of the Common Lisp and Standard Lisp options over the next few months.

1.2.3 POP

The POP-11 version of POP in the Sussex University POPLOG system will be the main provider of POP in the short term and is therefore the reference version. This needs complete documentation as a matter of urgency.

1.2.4 Mixed Language Working

It will be necessary to build systems that mix AI languages and conventional procedural languages. Currently POPLOG is a basis for mixed POP, PROLOG, LISP working and also allows linking in of other languages. It may be ported to new systems with a relatively small amount of effort (<1 MY) and it is therefore the preferred standard for mixed language working involving AI languages. POPLOG has a number of shortcomings however:

(i) The Lisp system is inadequate for anything beyond toy programs. The feasibilty of providing a realistic implementation is under investigation.

(ii) The Prolog system does not fully conform to the standard definition. Work is in progress to bring it into line.

(iii) POPLOG supports mixed languages by using a common low level virtual machine. This imposes some constraints that mean POPLOG does not achieve the same level of performance possible from single language compilers. For many applications the benefits gained by mixing languages will more than outweigh this loss but the possibility of further improving performance by modifications to the virtual machine will be studied.

POPLOG allows very close coupling of languages; the building of IKBS may also be expected to entail the composition of disparate systems which do not need such close connections. For this purpose Inter Process Communication (IPC) and Remote Procedure Call (RPC) standards are required and the IKBS programme will actively support activity in this area of the communications policy.

2. IKBS Research Tools

The following list contains those tools known to be of interest to several groups of researchers and will therefore be made more widely available through the IKBS infrastructure. This list is only provisional and will be revised and expanded in the near future.

2.1 Languages

APM — Logic programming
PS-Algol — Databases
OPS-5 — Rule based programming

3. Systems

The overall policy is to meet the majority of computing requirements through distributed systems of multi user and single user systems as set out in the Software Engineering policy (Appendix B).

3.1 Single User Systems

(1) Pending the evaluation of the current market no SUS will be purchased unless there is an urgent requirement that cannot be met in any other way.

(2) Evaluation of current products will allow the definition of supported and specialist systems by 4Q84.

3.2 Multi User Minis

(1) Multi user minis are the core of the academic infrastructure and will be the focus of investment for the next 12 months (i.e. to mid 85). Currently there are two supported systems for IKBS: GEC Series 63 and VAX compatible. There are currently 5 IKBS MUM academic infrastructure sites. Further expansion will be directly related to approval of major projects and is expected to include industry sites.

(2) The standard systems are expected to support up to 4 (VAX) or 8 (Series 63) simultaneous users of IKBS application software and in especially demanding cases the numbers will be lower.

(3) Beyond 2Q85 the upgrade path will be by the addition of networked single user systems. Each MUM will be expected to host up to 6 SUS.

(4) For research into IKBS virtual machines the High Level Hardware ORION is a specialist system.

3.3 Mainframe

(1) Most IKBS work is highly interactive in nature. A mainframe service for the most demanding users will therefore only be of benefit if the networking service is of a high bandwidth (greater than 4.8 Kbit/second) and the system response time over the network is low (less than 0.2 seconds). Providing such a service may therefore entail investment in high performance network links. Ideally the development of this service will include provision of the 5620 type of bit map displays, but current network protocols are particularly weak in the area of supporting highly interactive work.

(2) A mainframe service will also be provided in order to support small projects/pump priming usage in cases where new equipment is not justified and the MUM infrastructure is not appropriate.

(3) The most likely basis for this service over the next 12 months is the Amdahl UTS version of UNIX on IBM compatible equipment. However, resources will have to be deployed to make the IKBS languages available and this will have lower priority than MUM and SUS developments. A pilot service providing PROLOG and FranzLisp on the SERC RAL mainframes should be provided for 1Q85.

Appendix D

MMI Infrastructure Policy

1. Introduction

The MMI strategy calls for research to be undertaken under three main headings; displays, pattern analysis and the human interface. The computing infrastructure requirements of researchers working in these diverse fields are likely to be very different. This appendix is a general policy statement, defining the broad requirements, and identifying those issues that need to be resolved to produce a computing infrastructure that will meet the needs of a five year programme.

2. General MMI Policy Statements

The quality and degree of innovation of MMI research is critically dependent upon the nature of the available research tools. The primary objective of the Alvey MMI infrastructure programme is to provide Alvey MMI researchers with the best possible research tools. This may mean the purchase of a small number of the best research tools in the world (irrespective of country of manufacture) to ensure that UK researchers have the best possible chance of giving the UK a strong lead in this area. Ideally, such equipment should be accessible by any Alvey contractor.

It is desirable that there should be the maximum amount of interchange of tools and techniques between the various parts of the Alvey Programme. The MMI relationships with Software Engineering and IKBS are particularly strong. Therefore it is important that the standards and systems adopted in other parts of the Alvey Programme should be used by Alvey MMI workers. However, this policy will not be followed to the extent of restricting progress in MMI.

In many of the research areas encompassed by Alvey MMI, a significant amount of unco-ordinated effort in the past has been invested in tools for the collection and analysis of real-world data (e.g. speech waveforms, digital images, psycho-physical data, etc). This data has a life-time which is longer than that of current IT products. A national programme provides the opportunity for a concerted infrastructure activity that will prevent the duplication of activity, and facilitate the exchange of data, tools, code, results, etc. In order to protect the investment in infrastructure from continued technological change, the MMI infrastructure policy will place the primary emphasis on standards for data representation, and on software and hardware interfaces. The policy will also be reviewed at regular intervals to ensure that it evolves to meet the needs of both industrial and academic research.

In some aspects of the Human Interface Programme, it is particularly important to have the most advanced technology available, together with adequate reserves of computing power, so that the limits of human performance can be fully explored. These requirements are unlikely to be met by products designed for current commercial markets. After a detailed examination of these research requirements, the MMI Director will, where necessary, provide funding for the development of special purpose tools to meet these advanced requirements.

3. Standards and Systems

This section is a preliminary and incomplete statement of MMI infrastructure requirements.

3.1 Data

Wherever possible, standards for data collection and storage will be established, to facilitate sharing of data and the comparison of algorithms.

3.2 Operating Systems

Some IKBS activities would benefit from the provision of speech, image and human interface front ends. MMI workers will want to use IKBS techniques. Compatibility with UNIX is therefore to be preferred. However the message passing and synchronisation aspects of some versions of UNIX are not adequate for MMI work.

3.3 Software
Use of high level standards wherever possible, e.g. GKS.

Mixed language working is desirable to ensure the optimum use of code developed in different environments.

Establishment of algorithm libraries.

3.4 Hardware
Adoption of interface and highway standards to facilitate interconnection of devices and subsystems e.g. VME.

3.5 Systems
The provision of high performance work-stations is extremely important to most MMI work. There is a particular need for high resolution bit-mapped colour displays.

4. Implementation

A programme has already been initiated to examine the MMI infrastructure requirements in some detail, starting from these top level policy statements. The likely sequence of events will be:

1. Refine (and add to) top level statements.
2. Survey the MMI communities for their detailed requirements.
3. Identify those requirements that can be met from available equipment.
4. Propose specific infrastructure programme to meet the remaining requirements and to intercept future requirements.

Given the diverse nature of the MMI requirements, it is likely that full-time effort will be required to manage the MMI infrastructure programme.

Appendix E

VLSI & CAD Infrastructure Policy

1. Introduction

The purpose of this section is to indicate where the Alvey Programme in VLSI & CAD differs from the other enabling technologies with respect to its infrastructure requirements.

1.1 The VLSI technology programme (as opposed to the CAD programme) will have some needs for computing facilities, but these will normally be rather specialised to particular projects and linked to specialised hardware. One exception is the Device and Process Modelling Topic, which will benefit from the same infrastructure as the CAD programme.

1.2 In one sense, the Alvey Programmes in Software Engineering, MMI, and even IKBS, can be viewed as enabling technologies for the CAD for VLSI Programme. It might be hoped therefore, that the infrastructure requirements of these technologies would largely satisfy the CAD Programme. While this is partially true, the detailed requirements, so far as these are known at this time, are not identical.

1.3 Compared with the Software Engineering programme, the CAD Programme is more concerned with Application Software than with System Software. For example, participants are less likely to need to develop their own programming languages.

Conversely, they have other needs which are not so common in the other technologies. For example, in addition to the need for portability of software, they need to be able to transfer design data (representing an electronics system) between different design systems.

1.4 One significant practical difference between the CAD Programme and the other Programmes is the high level of existing *industrial* investment in software and hardware for CAD. This probably represents several times the amount which the Alvey programme can contribute in CAD. This is matched by a strong research base in industrial companies. The strength of the academic community and level of SERC investment is at least comparable to that in the other technologies.

There is also a large and growing market in CAD Systems. These products will substantially influence future trends in infrastructure as they have in the past.

2. CAD Interface Standardisation Project (CISP)

The CAD Programme has the benefit of a recently published comprehensive report (the CISP Report) which is highly relevant to the infrastructure requirements of the CAD Programme.

2.1 The purpose of CISP was 'to identify areas of weakness in generation and interfacing of electronic CAD software so that inter and intra company collaborative developments could be improved and third party software could be incorporated more easily.'

The report includes recommendations to collaborators on how to make the best use of existing software tools and standards, and to funders of collaborative projects on what programmes of development work should be initiated to allow future projects to be more efficient.

2.2 The study was carried out by 5 working groups covering the following topics:

1. Administrational and technical lessons from past collaborative projects.

2. Availability and development of CAD software engineering tools, high level programming languages and operating systems.

3/4. Instore and disc data representations.

5. Design representation and interchange of data.

6. Graphical data representation.

(The numbering of the study groups reflects the fact that groups 3 and 4 amalgamated.)

2.3 The main recommendations can be seen in the extract from the report attached to this appendix. References will be made later in this appendix, where appropriate, to specific recommendations in the report.

2.4 The report is the result of a 5 month study involving staff from 15 organisations on a part-time basis. This work was initiated before the Alvey Programme was in full swing. It was originally promoted by MOD with the intention of making the results available to the Alvey Directorate.

The report was only available in its final form a matter of days before this appendix was written. Therefore, although the Directorate has been represented throughout the study, it has not had time to consider the proposals in detail. The fact that some proposals from CISP are reported here should not be taken to indicate that the Directorate has adopted them. However, it is expected that many of the recommendations in the report will be strong contenders for forming a basis for the CAD Infrastructure Programme when details of that Programme are decided.

3. Standards Policy

3.1 Operating Systems

One important recommendation in the CISP Report is that it should not be necessary to agree a common operating system in order to collaborate in CAD software development. This is possible because a typical CAD program only interacts with the operating system in a fairly simple way (e.g. open file).

Instead they recommend (CISP Recommendations 3 & 4) that two machine and operating system independent interfaces be defined ('conceptual machine' and 'conceptual file') and implemented on a range of machines and operating systems. This recommendation is likely to be accepted by the Directorate, and would then form part of the Infrastructure Programme.

In practice, by far the largest majority of research, development, and production work in CAD for VLSI in the UK is done on VAX machines with VMS operating system. Although the report recognises that UNIX is the only existing common operating system, it does not recommend UNIX for CAD due to an alleged lack of file security and resilience.

The recommended interfaces, if implemented, would remove any controversy over machines and operating systems, but they cannot be recommended as standards at this time because they do not exist. Certain work has been done in this area and should be examined for suitability as a de-facto standard. (E.g. Kernigan's Software Tools in Pascal, ICL's conceptual machine, and RSRE's conceptual file.)

Pending the existence of such interfaces, no standard can be recommended in operating systems. Developers of new software are recommended to examine these existing interfaces, use them if possible, and construct their programs so that conversion to any emerging standard interface will be as easy as possible.

3.2 Programming Languages

The recommended language is Pascal (ISO Standard). However, it is recognised that certain important features are not defined in the standard (e.g. separate compilation and dynamic storage allocation). Extensions will be a subject for work in the Infrastructure programme.

Since so much existing software is written in Fortran, ISO Standard Fortran 77 will be provided on supported systems.

3.3 Graphics

GKS is recommended for graphics work. However, the report recognises that the full standard is not needed for most CAD work in VLSI, and that some graphical structures (e.g. nets) used in CAD for VLSI do not map efficiently onto the structures available in GKS.

3.4 Communications

The requirements for CAD for VLSI (electronic mail and file transfer) are expected to be met as part of the general Alvey Communications Policy.

4. Systems Policy

The majority of work in the short term will continue to be done on multi-user systems. The most common of these is VAX+VMS, and this will be a supported system. VAX+UNIX is already a supported system for the Software Engineering Programme, and is an approved system for CAD.

SERC has made a significant investment in Prime computers for the centrally supported CAD facility for academic use. For this reason, Prime (750 family) +Primos will be an approved system.

When academic researchers are included in a collaborative project with industry, they should adopt the standards being used within that project. When they are exceptionally working alone with an industrial 'uncle', they should adopt the SERC Common Base Policy.

It is too early to specify the supported single-user systems. It is expected that the systems chosen for the SERC Common Base will be acceptable to the CAD Programme.

5. Infrastructure Programme

It is too early to specify the Infrastructure programme in detail. However, several items have been identified above

which will form a focus for this programme. The CISP report is a rich source of recommendations. The aim of the Infrastructure Programme will be to define a Foundation upon which all future CAD research can be based to optimise the sharing of CAD tools, both those resulting from different research programmes and those arriving on the market from commercial vendors.

The infrastructure programme will probably take up to two years to complete. This will involve several items which are special to CAD. (E.g. data interchange standards, data modelling, etc.) After this time, these interfaces will be adopted as mandatory standards.

6. Recommendations

The CISP recommendations outline work which would be too ambitious for any of the companies in CISP to undertake on their own. Members of group X are very conscious of current world trends which are forcing even the largest companies to collaborate with each other and to seek national funding in order to provide resources for their future VLSI CAD programmes. We consider that given the large foreign investments already noted, there is a danger of UK "thinking too small". Further, in the particular area of CAD infrastructure, we consider it vital that these recommendations give rise to real projects and real products.

The recommendations set out below are presented in two sections. Section 6.1 comprises those tasks, from the recommendations of the CISP working groups, that group X (which includes the other group leaders) considers should be prime candidates for funding. Further recommendations to funders are contained in the group reports. Section 6.2 identifies other tasks which group X itself considers are as important as those of Section 6.1 and which have not been highlighted by the working groups. Both sections are "recommendations to funders".

It is expected that those seeking technical guidance from the CISP reports will study the working group reports, which contain recommendations of good technical practice.

6.1 Group Recommendations to Funders

The following recommendations are given in working group order, not in order of priority. However they do fall naturally into four subject areas.

6.1.1 Communications

CISP Recommendation 1: Data Links between Collaborators

PSS X-25 data links should be provided between collaborating companies and their provision should be made an allowable cost under the Alvey scheme.

Group 1 Recommendation 3

6.1.2 Provisions for CAD Tool Portability

CISP Recommendation 2: Consolidation of Pascal

Pascal is recommended as the current preferred language for CAD tool development. A study must be undertaken to review existing compilers, to define required extensions to ISO Pascal, and to define how these extensions can be mapped onto existing compilers. This study must be followed up immediately by software development to produce dialect-converting and source-auditing support tools. In the longer term, the interface routines developed for use with Pascal should be reviewed in the context of Ada.

Group 2 Recommendations
7.2.0.1, 7.2.0.2, 7.2.0.3, 7.2.0.d
Group 1 Recommendation 1

CISP Recommendation 3: Conceptual Machine

Set up a project to produce an operating system independent Conceptual Machine interface. The binding of this interface to all of the common high level languages, including Pascal and ALGOL68, will be defined. The project will implement and support the interface on a range of commonly used computers. The Conceptual Machine interface provides a machine-independent foundation upon which portable software can be built.

Group 3/4 Recommendation 12
Group 2 Recommendation 7.2.0.c
Annexe 1

CISP Recommendation 4: Conceptual File Manager

Set up a project to recommend and to develop facilities, at the Conceptual File Manager level, to be built on top of the Conceptual Machine. This project would define powerful and portable means of carrying out many of the functions usually offered (nonportably) by operating system file handling facilities.

Group 2 Recommendations
7.2.0.1, 4.3.1
Group 3/4 Recommendation 13.
Annexe 2

6.1.3 Data Modelling and Interfacing Techniques

CISP Recommendation 5: VLSI CAD data models

Identify desirable software tools for the support of data modelling work. Determine the extent to which a common set of data models may be produced and create data models to cover as many aspects of VLSI CAD data as possible. This may involve a short research study on the data modelling of behavioural data. Review existing design description language to identify how completely they capture the types of data within the scope of the data models.

Group 5 Recommendations
Annexe 3
Annexe 4

CISP Recommendation 6: Interface Techniques

Identify or produce a suitable parser generator to assist collaborating companies in the generation of translators and interface formats consistent with a data model they have agreed between them.

Group 5 Recommendations

CISP Recommendation 7: Mask Transfer Interchange Format

Carry out a study of possible extensions to the CIF interchange format. The aim is to define a CIF superset forming a more powerful standard interchange format for mask data.

Group 6 Section 6.2

6.1.4 Graphics Support

CISP Recommendation 8: GKS Applied to CAD

Conduct a study of GKS applied to CAD and generate (starting from a subset of GKS) a standard of more relevance to current and future CAD work.

Group 6 Section 6.2

6.2 Additional Recommendations to Funders

Recommendation 9: Network Management

Group X recommends that work is funded on network strategy for CAD systems, and on development of a software tool kit to support such strategy. In particular, consideration should be given to what facilities, above the level of file transfer, would be needed. Work should start from the premise that the VLSI development environment of the immediate future will often have UNIX-based CAD workstations communicating with VLSI data management systems which do not run under UNIX, notably those running under ICL VME and VAX VMS. See the group 2 report for details.

CISP Recommendation 10: Configuration Management of Distributed Data

No working group was specifically charged with the study of change control requirements in a distributed network environment. Several groups made reference in their reports to the importance of this topic and its central place in CAD infrastructure. Group X recommends that the Alvey Directorate should commission a study contract to investigate requirements for configuration management of distributed data. In particular, the study should recommend what network facilities, Conceptual Machine and Conceptual File facilities need to be provided to support safe updating of design data on a distributed database. Both group 2 and group 3/4 Recommendation 15 touch upon this topic.

When the Study Contract is complete, the Directorate should stimulate development of the techniques and support tools that are shown to be needed.

CISP Recommendation 11: Development of Standard Man-Machine Interface Methods

Man-machine interface (MMI) are referred to in the group 1 report, and their graphical interaction requirements are covered in the group 6 report. It has not been possible within CISP to explore this subject fully, but it is acknowledged that high-level standardisation of MMI within a project is important. It is recommended that a conceptual model and a methodology should be defined for the specification of man-machine interfaces.

CISP Recommendation 12: Training for Existing Staff in New Techniques

To maximise the effectiveness of the new tools and techniques, which are funded as a result of this report, the Alvey Directorate should consider whether associated training and retraining costs could be made allowable costs under the Alvey scheme. For instance, "conversion course" retraining for staff skilled in the use of FORTRAN would help them to make the most effective use of Pascal. Similarly wider appreciation of the "layered-software" approach to building CAD data management systems would be of value.

CISP Recommendation 13: Training the New Generation of Engineers

The Alvey Directorate should encourage and assist the establishment of university training in data modelling techniques.

Annexe (to Appendix E) Membership of Group X

The reports of the individual working groups have been reviewed and this summary report prepared by a group referred to as "group X". The membership of this group is as follows:

Mr. H.G. Adshead ICL
Mr. K.R. Bennett STL
(Chairman of group X)
Mrs. A. Carter GEC
(Chairman of group 6)
Mr. B. Cosgrove ICL
Mr. E.J. Dowling FCSL
(Chairman of group 2)
Mr. P.J. Horth BT
(Chairman of group 3/4)
Mr. J.C. Howes FEL
Mr. E. Knight ICL
(Chairman of group 5)
Mr. S.P. Lane STL
(Chairman of group 1)
Dr. N. Peeling RSRE
Mr. F.R. Ramsey FEL
Mr. R.A. Reynolds BT
Dr. T.L. Thorp RSRE
(Technical sponsor)
Dr. J.D. Wilcock Plessey-Caswell

Group X meetings were also attended by:

Mr. M.J. Newman Representing the Alvey Directorate, CAD for VLSI.