Computing at CERN
in the
LEP Era

May 1983

COMPUTING AT CERN IN THE LEP ERA

CONTENTS

29 April 1983

The introduction of the large electron-positron storage ring LEP into the CERN programme has provided an opportunity to consider the computing and communication facilities which will be required for LEP experiments and the whole CERN physics programme at the end of the 1980's. The overall size and complexity of the detectors, the large amount of data which will be produced and the large distances between the experiments and the CERN site have shown the need to have first-class facilities which will enable the physicist to participate fully in the experiment.

The recent discovery of the Intermediate Vector Boson at CERN has highlighted the importance of being able to analyse large amounts of data from a well-designed and calibrated detector in a short time scale. A similar situation will surely exist in the early phase of LEP experiments.

The recommendations given here by the LEP Computer Planning Group (and submitted to the Directorate) clearly show that we are living in a period of very rapid developments in the field of electronics, computing and communications. This is the first time that such a comprehensive study has been carried out involving CERN research and accelerator staff together with users and outside experts.

I would strongly recommend that a workshop be held every $\sim$ two years bringing together the same people to update the recommendations in light of present developments or new information.

Erwin Gabathuler

Erwin Gabathuler

Report of the LEP Computing Planning Group

May 2, 1983

---

1. INTRODUCTION

The LEP Computing Planning Group met for the first time on April 24 1982 and held a total of 17 meetings. The membership is listed in Appendix A. Working groups were set up to study LEP computing needs in three areas:- (WG1) Data acquisition and monitoring (WG2) Networks and Telecommunications (WG3) General computing services and facilities. All of the working groups set up some smaller bodies to study particular topics. WG3 held an open meeting to solicit user input on 22 September 1982 and WG2 and WG3 jointly held a meeting on 1 December 1982 to invite comment on their draft reports, which had been distributed to outside experts. All three groups have now produced extensive reports, which contain much background analysis and which make many detailed recommendations. These working group reports are included as Appendices B,C and D. The summary of the report of the Microprocessor Working Groups, which had been set up by the CERN Directorate independently from the present studies, is attached, for completeness, as Appendix E. The full report (about 60 pages) is available from Mrs Avis Hoekemeijer/DD.

The first Chapters of this report give a very brief description of the role and style of computing in the coming decade, with special reference to the needs of the LEP experiments. In Chapter 7 we highlight, in Sections 7.1-7.4, those technical recommendations that we think will be especially important for computing in the LEP era, and in Section 7.5 we make some proposals for appropriate organisational structures within which the work might be performed. In Chapter 8 we summarize the requirements for staffing and money that will be needed to carry out our recommendations, and make some comments on the total requirements in the context of the present CERN budget.

Readers who would like to understand the overall background to our recommendations are encouraged to read the complete Working Group reports. For the background to our individual recommendations, a set of references to the Working Group reports can be found in Chapter 9.

This report is the first for a long time to attempt to cover the place of computing and allied techniques across the whole CERN programme. It does not, however, give a complete coverage: its main aims were to make proposals for new capital equipment and technical facilities that will be needed by the forthcoming experimental programme. It only touches on the computing needs of the CERN accelerators, including LEP. It also largely ignores classical fast electronics, special trigger processors, Camac, and administrative and office computing. All these activities have staff and budgets in various Divisions which we have not included in our estimates, but which will have to be taken into account by the CERN Management when fitting our proposals into the overall CERN programme.

2. THE IMPACT OF COMPUTING IN PARTICLE PHYSICS

Ever since experiments began at CERN computing has been playing a growing role in particle physics. Early on (about 1965) the emphasis was on the new scientific mainframe computers. CERN acquired large processors (CDC 6600 and then CDC 7600) and used them mainly for batch processing. During the 1970s the changing technology led to the rise in importance of the minicomputer for accelerator control and for data acquisition at the experiments. Today CERN has at least 250 minicomputers being used in this sort of environment. At the same time the emphasis at the computer centre was changing, with the performance and reliability of peripherals becoming at least as important as the raw central processor power, and with punched card input being superseded by access from terminals spread all over the site.

During the past few years technology has again pushed new products to the fore. Microcomputers now offer the optimal solution in many environments; 32-bit minicomputers offer significant power and software for many general-purpose applications; and new super-computers continue to offer batch processing at progressively lower prices. All of these technologies will have their place in a modern particle physics laboratory like CERN. Indeed we can already see the next technology, namely the custom design of special purpose Very Large Scale Integration chips for physics, starting to be considered for some applications.

In parallel, the need to interconnect the growing number of computers and to offer access to computers from remote locations on the site, led to the development of various types of network. Although this sort of requirement was something of a CERN speciality in the 1970s, nearly all commercial computer systems available today lay heavy emphasis on the ease with which they can be interconnected through networks.

During the late 1970s expenditure on physics computing at CERN has been running at about 7% of the total materials budget (roughly 22 Msf out of 300 Msf). Of this some 18 Msf was typically spent through the Data Handling (DD) Division budget and 4 Msf was spent by the Experimental Physics (EP) and Experimental Physics Facilities (EF) Divisions, with a total of about 10 Msf per year being committed for new acquisitions (purchase or rental). The corresponding staff figures are about 200 staff working in DD Division and a large number of staff, difficult to determine precisely, more directly associated with experiments in EP and EF Divisions.

A graph showing the steady rise in central processing power available and used at the CERN computer centre can be found in Appendix D, Figure 5. The steadily decreasing cost of computing equipment has allowed CERN to provide, at roughly constant cost, an order of magnitude increase in computing power over the past decade.

3. COMPUTING AT LEP - GENERAL POINTS

The LEP experiments will give rise to a number of new computing problems because of the way in which they differ from the type of experiments which have been common in particle physics until now. Experiments at the $p\bar{p}$ collider already show some of these characteristics.

3.1. Number of Collaborators

LEP experiments will be very large, some having more than 200 physicists and up to 300 collaborators if we include support staff. Communication among these staff, and the possibility for them to connect to their home computers from CERN and to the experiment from their home institute, will become crucial in determining the efficiency with which the experiment can be carried out. We note in passing that the presence of multiple PTT administrations in Europe means that the job of ensuring good communication among the members of a large team is much more complex here than in the USA.

3.2. Physical Size

The nature of physics at LEP has the consequence that the size and complexity of the detector increases and this then feeds through into:- greater complexity of the read-out electronics; a greater need for online monitoring; extra requirements for offline batch capacity; and a requirement to improve the front-end processing (triggering and data reduction) in order not to overload the later data analysis stages.

3.3. Logical Complexity

Another consideration is that there will be a very large number (say 100-200) of programmable processors in each experiment. The principal effect of this will be to increase the complexity of the experiment, and collaborations would be well advised to look very carefully at the related systems design.

Some aspects of the experiment that will probably be amenable to better design are:-

- Configuring the computer(s) and microprocessors that control a detector sub-system in such a way that the whole assembly can be physically moved to and from a test beam location and still be able to run with the same programs as in the normal environment of the experiment.
- Designing the dialogue between the detector sub-systems so that control- and data-flow are handled correctly when the whole detector is assembled together.
- Arranging to present the mass of data generated by the experiment in such a way that the physicists responsible for a run are neither flooded by irrelevant information nor deprived of access to crucial details of detector performance.

4. THE NON-LEP PROGRAM

Although this report mainly concentrates on the computing requirements of the LEP program, we need to remember that CERN will have to carry out an active program of fixed target and $p\bar{p}$ collider physics throughout the LEP construction period and beyond. We have taken this program into account when preparing our estimates of future requirements. Indeed there is a distinct possibility that during the next few years the increasing requirements for batch processing capacity and interactive graphics arising from the $p\bar{p}$ collider program will overload the existing facilities, independently of the load resulting from the preparations for the LEP experiments themselves.

We have assumed that developments carried out with LEP experiments in mind will form an adequate basis for dealing with any other non-LEP experiments that may be started in the meantime.

We should also emphasise that we have attempted to treat the needs of the accelerator constructors and the designers of physics equipment on the same footing as those of the physicists who will eventually use LEP. Many of the requirements of these groups for computing services, which have had a very different profile in the past, are now starting to look increasingly similar. Examples include the need for interaction in general, for databases, for computer aided design, and for microprocessor support. We have noted, without allowing for it in our proposals, that theoretical physicists have started using considerable amounts of computing power for lattice gauge computations.

5. A MODEL FOR LEP COMPUTING

It is possible to merge the information given in the Working Group reports to obtain a picture of the most likely scenario for LEP computing in the 1987-1988 timescale. At each experiment there will be several 32-bit minicomputers for the overall data acquistion and control, driving many (up to several hundred) microcomputers. The data from the detectors will arrive through Fastbus and Camac; terminals will be connected to computers inside the experiment via a Local Area Network; and there will be reliable easy-to-use connections from the experiment to the computer centre.

At the same time it will be possible to have access to the experiment from outside CERN and vice versa. The networking inside CERN will be based on a new backbone network and gateway connections to a number of special purpose networks will be available. The outside networking will be based on the X25 public packet switching services.

Although good batch services will be needed, especially to cope with the high volume of on-site analysis that is predicted for the first few years of LEP physics, the emphasis will be moving more towards easy interactive use of computing facilities by physicists and programmers. High performance graphics workstations will have come into their own as tools for the design of detectors and complex electronics, for program optimisation and for event scanning.

6. DECENTRALISED COMPUTING - HOW TO AVOID CONFUSION

It is clear that the cost of computer hardware and the associated electronics is still falling extremely quickly. This continuing trend permits both industry and CERN itself to develop new approaches to old problems with disconcerting speed. However the integration of hardware and software from different manufacturers into a coherent piece of equipment which solves the physicist's or engineer's problem is as hard as ever to achieve.

There is a real risk that an organisation like CERN will find that it is wasting significant effort in this integration. The only way to control this is for the laboratory to encourage the adoption of **STANDARDS.** Where possible these should be internationally agreed standards, but if these do not exist then CERN should select appropriate national, de facto industrial, or even as a last resort in-house, standards for use by all staff.

The encouragement to use the selected standards should be that central support (whether hardware advice, repair services, software services etc.) should not be made available for **NON-RECOMMENDED** systems. Any requests to connect non-recommended systems to CERN facilities should be decided on a case-by-case basis, and any connections that are approved should be required to conform to interface specifications laid down by CERN.

We further believe that some central control of computer purchases is desirable, while noting that electronic components are now so cheap that each and every CERN engineer and technician has the personal responsibility for whether or not he or she follows general guidelines in this area. The implications of the high and frequently dominating cost of software in development projects must be brought home to people who wish to go it alone.

7. CONCLUSIONS AND RECOMMENDATIONS

7.1. Data Communications

All of the LEP groups have emphasised that if they are to function as true 'collaborations' then it will be vital that communications facilities between CERN and the outside institutes involved in the CERN program, as well as those inside CERN, are reliable and easy to use. We must aim for a situation where the physicist or engineer can contribute almost as effectively to the collaboration whether working at CERN or at the home institute.

Furthermore, in the light of experience it has become clear that CERN needs an integrated communications system which covers the whole range of computer applications.

These general requirements lead us to make the following related set of Recommendations 1-7.

7.1.1. General Network Architecture

Recommendation 1
We recommend that the architectural model that CERN should adopt for its telecommunications planning for the next decade should have a high-speed general-purpose backbone network covering the whole of the extended CERN site, with gateway connections to certain special-purpose networks.

The special-purpose networks will include the future LEP control network, the PS and SPS control networks, Local Area Networks installed at LEP experiments or to serve CERN buildings, and the external networks. When a new CERN digital telephone exchange is installed then it will (logically speaking) become another of these special-purpose networks. If it becomes desirable to install a special high performance network inside the computer centre then that would be yet another such network.

Requirements for very high speed (8 Megabits per second and greater), large volume (several Gigabytes) point to point transmission should be satisfied by providing dedicated links, using the same technology and standards, but independent of the backbone.

This recommendation on overall architecture should be studied by a small group responsible for making a more detailed assessment of the choices that are open. Definite decisions should have been reached by the end of 1984 so that the special purpose networks can proceed with well-understood interface specifications.

The effort required for this study will total about one manyear per year in 1983 and 1984.

7.1.2. Basic Communications Infrastructure

Whereas in the past data communications systems in CERN have used different types of cables and basic infrastructure, there will be considerable advantages in standardisation at this level.

Recommendation 2
We recommend that, where appropriate, transmission and multiplexing of the network connections and communications links over long distances should follow the CCITT G-700 Series specifications for PTT systems.

During 1983 a test installation between CERN Meyrin and CERN Prevessin has been funded by the LEP budget. The LEP project planning includes allowance for this infrastructure for the accelerator controls. For the physics infrastructure 100 Ksf per year should be budgeted in 1984 and 1985. At present 2 staff are working in this area, and 2 extra will be required, 1 for the accelerator aspects and 1 for the physics side, starting in 1984.

7.1.3. Standardisation of High Level Protocols

We now have enough experience of networks to know that the provision of a path between two computers, or from a terminal to a computer, is the start and not the end of the user's problems. For the user to be able to do anything useful, such as transferring files, reading electronic mail, or working on the remote computer, it is necessary for both parties to agree on the details (the high level protocols) of how they will speak to each other.

Even computer manufacturers have taken a long time to define these high level protocols and to implement them successfully on equipment that is completely under their own control. The situation that occurs when computers and terminals from different manufacturers are linked together is an order of magnitude more complex, and CERN is one of the sites that has had to tackle this problem, not without success. Various international standardisation efforts are under way, but progress is slow, and in this area products conforming to any new standard will not become available in the near future.

Until the time when international standards are fully defined and supported, CERN and the European particle physics community will need to decide which de facto or other standards and protocols they wish to recommend and provide support for. Since, even in the medium term (5-7 years), we will have to live with the present multiplicity of protocols, we believe that the only way of implementing successful access to services across networks (interworking) will depend on translation between protocols. We must point out that these problems of protocols and interworking are probably even more severe for the external connections than for those on the CERN site, there being more technical constraints that are outside the user's direct control. Very valuable work is being carried out by an ECFA working group in this area (see next section). As far as possible, common solutions should be implemented for off-site and on-site connection problems.

> Recommendation 3
> We recommend that CERN, in collaboration with the European particle physics community, should adopt and support a set of higher level network protocols.

The staff effort required has been incorporated into the following sections (7.1.4 to 7.1.6).

7.1.4. External Particle Physics Network

> Recommendation 4
> We recommend that the present limited effort to provide basic access across public packet-switched networks should be expanded in order to provide reliable terminal access, file, mail and message transfer for all particle physics institutes collaborating in the CERN program. Technical collaboration in this area should follow the guidelines established in the report ECFA/82/60. If it is to be successful this proposal will require the active support of a number of external laboratories.

A limited X25 service will start in May 1983, using a Modcomp computer of the type used in CERNET as a switch and gateway between CERNET and the external public networks. CERN expenditure at the level of 100 Ksf in 1983, 200 Ksf in 1984 and 300 Ksf in 1985 should be foreseen to cope with the growth expected in this service. Users will be billed for their real "connect-time" use of the external services and, at least during the initial period, DD Division will pay for the "standing charge" (abonnement). We foresee that in the longer term a line item of 100 Ksf per year should be included in the budget to cover computer and X25 equipment maintenance, and for the "standing charge". At present the equivalent of 2 manyears per year are committed to this work and to ECFA protocol studies. This needs to rise to 4 manyears per year by the end of this year and reach 6 manyears per year in 1985, to cover both the X25 service and the CERN share of the work on protocol conversion.

Corresponding investments will be have to be made by several countries or institutes if this external network is to be successful.

7.1.5. CERN Backbone Network

We should aim for the new backbone network to have a capacity of at least an order of magnitude more than that of CERNET, which is currently acting as the backbone.

> Recommendation 5
> We recommend that the future backbone network should, wherever possible, incorporate standard commercial products, and efforts should begin immediately to identify such products, with a view to having the network partially operational from the end of 1985. The integration of this backbone with the LEP controls network and with the other communications facilities on the CERN site, including Index, the present CERNET, and a possible future CERN digital telephone exchange, is of crucial importance. Until this new backbone can be installed the capabilities of CERNET should be monitored and any improvements kept to the essential minimum.

Estimating money and manpower required for this project is not easy since the availability of partial or complete solutions from industry is not obvious. For comparison, CERNET required about 3 Msf and 60 staff-years of initial development. The total expenditure required for the backbone is estimated to be 2.2 Msf (of which 200 Ksf are for gateway connections), over the period 1983/1988. Reasonable build-up would give expenditure of 100 Ksf in 1983, 200 Ksf in 1984 and 500 Ksf in 1985. When installation has been completed the expenditure for maintenance and operation should stabilise at about 300 Ksf per year. The present effort devoted to CERNET, which is the nearest equivalent to the backbone, and gateway development, is the equivalent of 8 manyears per year. The total effort required to implement the backbone and the first gateways, while keeping the 100 CERNET connections working, is likely to peak in 1985/1986 at about 12 manyears per year.

7.1.6. Local Area Networks

Local Area Networks will start to become widely available during 1983. Provided that their reliability, performance and price are all satisfactory they will be an attractive solution for linking together terminals and computers in a limited geographical area such as one building or an experiment.

As discussed in Section 7.1.1 on General Network Architecture, we expect these Local Area Networks to be connected to the general-purpose backbone network in order to access other computing facilities at CERN and outside. Each connection of a new type of Local Area Network will absorb significant resources, not only for the lower level hardware but also in order to provide the correct high level protocols.

Several standardisation efforts are being undertaken in this field. At the lower levels Ethernet has been specified for some time by a consortium of manufacturers, and interface chips for the low level protocols are now becoming available. Also at the lower levels the IEEE 802 Committee is preparing a draft standard that will include Ethernet and a number of alternative (and incompatible) technologies. The European Computer Manufacturers Association (ECMA) have agreed a proposal for software protocols that will provide task-to-task communication based on underlying Ethernet hardware.

> Recommendation 6
> We recommend that CERN should choose and support a standard Local Area Network both for use in LEP experiments and for general purpose applications. There is one obvious candidate at present, namely Ethernet.
>
> Any connections to the CERN networks of other Local Area Networks, which might be introduced onto the site as part of special systems, should be required to conform to CERN interface specifications. Such networks will not normally be supported by CERN, and their connection should be subject to approval by the leader of the Data Handling Division on a case by case basis.

DD Division has already been requested to obtain and evaluate an Ethernet system, and to provide the basic hardware and software expertise needed to make Ethernet usable by others at CERN. Funding of 100 Ksf has already been found for 1983, but there will be a need for 200 Ksf per year in 1984 and 1985.

It is assumed that users will pay for their own hardware when they decide to use a Local Area Network. The number of staff needed for support activities will depend both on the number of Local Area Networks installed and on CERN's success in limiting the support to a single recommended type. Assuming an introductory phase where only a single recommended type is involved, 2 manyears per year are needed now, rising to 3 manyears per year in 1985.

7.1.7. Terminal Connections

There are two general approaches by which terminals can be connected to different computers. One is circuit-switching, in which a physical path is created between the terminal and the computer. The Index system, widely used at CERN, is an example of this approach, and the new generation of digital telephone exchanges will allow for circuit switching of terminals as well as of voice. The other method is packet-switching, in which, instead of a physical connection based on cables, a logical connection based on a data communications network is created.

In the future we can expect both methods to evolve, and there will be some areas where each approach has strong advantages over the other. There is no sign that over the next decade either technology will eliminate the other, whether speaking in the CERN or world-wide context. We expect, however, that the present predominance of circuit-switched connections at CERN will diminish when cheap intelligent connections to networks become available.

We can see little hope in an organisation as complex as CERN of avoiding the need to support both types of terminal connection. It is our hope that the General Network Architecture, see Section 7.1.1, and CERN Backbone Network, see Section 7.1.5, will be sufficiently well designed to protect the user from any related problems.

> Recommendation 7
> We recommend that staff working on the General Network Architecture and CERN Backbone Network must aim for good integration of terminal connections whether they are made through circuit switched or packet switched techniques.

The recommendation does not carry any directly related costs, since we expect users to continue to pay, as now, for their terminal connections. However, to emphasise the importance of investment in this area we would like to point out that CERN has over 1500 terminals and that there are about 15 staff whose primary activity is related to installing these terminals and to keeping them working. The current backlog for installation of terminals by the DD Terminal Pool is just over 250. Extra manpower for maintenance, possibly from outside contractors, will be needed as the number and complexity of terminals grows.

7.2. Data Acquisition

7.2.1. Fastbus

Fastbus offers at least an order of magnitude increase in speed over Camac and it has an architecture which permits the synthesis and easy modification of complex multi-processor data acquisition and trigger systems. For these reasons we expect that it will be widely adopted by the LEP experiments. Experience with Camac has shown that significant economies are possible when central support of general purpose hardware and software are available.

> Recommendation 8
> We recommend that development and maintenance of general purpose Fastbus hardware and software should be provided by CERN. Furthermore, we recommend that CERN should encourage involvement of European industry in Fastbus, in order to ensure that Fastbus components are available from a range of suppliers.

The LEP experiments will, of course, pay for the general purpose Fastbus modules that they need (either as direct purchases or as Electronic Pool allocations). Start-up funding of the development and maintenance services for general purpose Fastbus hardware will cost 200 Ksf in 1983 and 400 Ksf per year in 1984/1986. Following this start-up phase the budget requirement will stabilise at 200 Ksf per year, corresponding to a production phase in which tried and tested modules are entirely paid for by the end-users.

The equivalent of 10 manyears per year is being invested in Fastbus development today (about 8 for hardware and 2 for software), without counting any workshop effort. This needs to rise to a total of about 18 manyears per year for development (13 for hardware and 5 for software) during the years 1984 and 1985. In order to reach this level there will be a need for 4 new hardware posts and 2 new software posts.

Assuming that a significant amount of the maintenance can be contracted to outside industry, as is customary for present electronics, an extra 4 manyears per year of technical support will be required in the Electronics Pool (present strength 25) on a long-term basis. Some of this support effort should become available by the start of 1984 at the latest.

7.2.2. Data Acquisition Computers

CERN's continued policy is to support minicomputers from both Norsk Data and Digital Equipment. However, taking into account the choice of computers that have been acquired by European physics institutes, and the indications of the LEP collaborations themselves, we recognise that the VAX family of computers will be the preferred 32-bit computers for LEP data acquisition systems.

Although the LEP experiments will only start to take data in about 1988 the choice of data acquisition computers is already fairly urgent since detector tests will start soon. We believe that those collaborations that aim to use a single minicomputer family for all of their test set-ups are taking a sensible decision.

As outlined in Chapter 3 it is easy to see that there may be many different types of processor involved in data acquisition systems for LEP experiments. This may include special trigger processors, commercially available 8- and 16-bit microprocessors, computers (probably 32-bit) for the overall control, and workstations for event display. These processors may be linked together through a variety of different bus and link systems, including Camac, Fastbus, VME, Local Area Networks, and CERNET and the new backbone network.

Integration of these different elements, each of which may provide the optimum solution to a given problem, will not be an easy job. The requirement to be able to easily relocate sub-detectors to test beam locations without major changes to the computing systems will pose further problems. Successful construction of data acquisition systems for LEP experiments is therefore likely to depend on correct selection and support for standard hardware and software components by the collaborations and the associated CERN support teams. Collaboration between the many people involved will be important during the whole construction phase, and some of the organisational issues are discussed in Section 7.5.3.

We assume that the collaborations will themselves acquire the data acquisition computers for their experiments. CERN itself will need some computers in order to carry out its share of the hardware and software development and support, and this will cost 100 Ksf per year. As a longer term action, existing test systems in the Online Computer Pool (50 HPs, Caviar..) which are no longer capable of running the latest software, or which are becoming impracticable to maintain because of age, must be replaced. A sum of about 400 Ksf per year will be needed in order to replace between 10 and 20 test systems annually. In order to cope with the increasing number of widely distributed computers the present Online Computer Pool team of 10 should be increased by 2 extra posts.

At present there are about 15 staff working directly on the development and support of software for all data acquisition systems. Especially taking into account the work that will be needed to ensure smooth operation of the range of new communications facilities, and assuming that support effort for existing data acquisition systems on 16-bit computers can be brought to a minimal level during 1983, 4 extra programmers will be needed in this area by 1984. Further staff would be needed if more than one data acquisition system has to be developed and supported for LEP experiments.

7.2.3. Guidelines for Electronics Systems Design

It is important, in order to economise on staff and money, that electronics designed for LEP should conform to two guiding principles. The first is always to follow recognised standards, and to avoid, for whatever reason appears good at the time, excessive optimisation for any individual application. The second is to incorporate features for self-diagnosis and fault-finding into all modules, in view of the high cost of using people to diagnose faulty electronics in a complex environment. The second point will be critical for the LEP experiments, since they will not want to waste precious initial beam time while bringing their electronics to full operational efficiency.

7.2.4. Support Services

Even though the designers will have done their best and produced highly reliable modules, with many self-test facilities, we have to recognise that the environment of an experiment is complex and during data taking there is always intense pressure to repair any malfunctioning computer or electronics as quickly as possible. In those circumstances there is no substitute for a good stock of tested spares and for access to staff who are experienced in the diagnosis and repair of these faults.

> Recommendation 9
> We recommend that CERN should continue to make adequate assistance available to experiments by suitably strengthening the CERNET support, Electronics Pool and Online Computer Pool teams, in order to handle the new load from LEP experiments.

Suitable strengthening of the hardware support teams to cope with the extra load coming from LEP experiments has been covered in the relevant sections (All of 7.1 for CERNET; 7.2.1 for the Electronics Pool support for Fastbus; 7.2.2 for Online Computer Pool support for data acquisition computers; 7.3.7 for Microprocessor Support).

7.3. Centrally organised facilities

7.3.1. "Private" Computer Centres

For the past decade CERN has operated the "one-third / two-thirds" rule, which requires that only one-third of the computing for any experiment is carried out at the CERN central computers, and that the other two-thirds of the total computing load have to be handled "at home" on computers in universities or institutes to which the collaboration has access. The motivation for the rule is at least twofold, principally to ensure that physicists are able to return "home" and contribute to the life of their "home" institute when data taking is complete, and secondarily to limit somewhat the capacity needed at CERN.

There has traditionally been some indulgence shown during the start-up phase of new experiments when calculating compliance with the rule. It is clear that if several experiments start up simultaneously on a collider and find interesting physics within the first few months then the pressure to do all of the initial computing at CERN will be very strong. Several of the LEP experiments have indicated that they might like to install their "private" computer centres to cope with this situation. There is a significant risk that the sociological objective of the "one-third / two-thirds" rule will not be attained if "private" computer centres of this type are installed at CERN.

If the collaborations pay for the infrastructure of their computer centres (building, power supply etc.) and look after the operation of the computers (operators, training, user support, paper supplies, tape purchasing etc.) themselves and do not need any connection to other computer facilities (Index, CERNET and their successors) then they will indeed be "private" and it would be reasonable that only collaborators should have access. We do not think that, in practice, this approach can work.

It is important to realise that there are important manpower implications if the computer centre is asked to run a "private" computer. Depending on exactly how the machine is run (as an integrated part of the CERN centre, or completely stand-alone) between 2 and 5 staff will be required for general operation and support, and 5 people, who could be contract staff, will be needed for 24-hour per day operation of the peripherals, especially magnetic tapes.

Many of the initial discussions have involved "private" computer centres that do not have a well-balanced set of peripherals for general "public" operation. They might therefore involve heavy CERN investment if they are to be used partially to augment the "public" capacity. In addition the presence of such "private" machines, if they are supposed to partially contribute to "public" capacity, is likely to severely limit the future freedom to optimise computer centre configurations.

> Recommendation 10
> We believe that the installation of very large "private" computers at CERN are likely to consume significant CERN resources. Such requests should be closely analysed at Directorate level, in the light of an overall CERN policy.

7.3.2. Central Batch Capacity

It is clear that increased central batch capacity will be required, both at CERN and in the other computer centres used by the particle physics community, by 1987. Assuming that the "one-third / two-thirds" rule is rigidly applied, an increase of a factor 2 by 1987 might just be sufficient, though reasonable arguments have been presented to us suggesting that in reality the factor might turn out to be as high as 4. Pressure to upgrade the capacity at an early date might arise as a result of $p\bar{p}$ collider physics results.

> Recommendation 11
> We recommend that CERN should plan at least to double its central batch capacity by 1987.

The cost of a factor 2 increase in capacity, if it were needed today, would be about 11 Msf and payments could be spread over several years. Support for this batch capacity increase would not require any change in numbers of CERN staff, since it is assumed that the numbers of mainframes in the centre would not change, and the services could be optimised accordingly. However, if magnetic tape handling on the computer centre mainframes continues to rise at the present rates, then the number of service contract staff working in this area will have to be increased. In that case an expenditure of 500 Ksf per year to cover the extra staff will be required by about 1986.

7.3.3. Production capacity based on emulators

Money and manpower are already committed for the development of 3081/E emulators, principally with a view to their online use in high level triggers. With some additional investment 3081/Es could also provide a useful amount of production capacity for offline batch processing. At least one LEP collaboration is interested in the possibility of using a bank of 3081/Es as a "private" production facility.

It is important to realise that while a bank of emulators could **add** significantly to the production capacity available in the computer centre, they are not an **alternative** to the major mainframes, on which they depend for services such as compilation, debugging and file handling. Furthermore emulators cannot contribute to the capacity required to handle the day-time (interactive) load in the computer centre.

Useful experience with the use of emulators for production has been gained over the past 18 months by running a bank of 4 to 5 168/Es attached to the computer centre through CERNET. The main problems that were observed concerned the non-negligible fraction of the central resources that was required to support the very heavy CERNET traffic, and the serious impact that was sometimes felt by other computer centre users when the emulators pre-empted peripherals such as tapes, disks or the mass storage system (MSS). In addition it has proved difficult, for a variety of operational reasons, to use the emulators as intensively as a normal batch service, and the long-term duty factor has been about 50%.

The 3081/E will be more powerful than the 168/E, and easier to use. The favourite model for a bank of 3081/E emulators would avoid the problems with computer centre resources by attaching 4 or 5 emulators directly to a dedicated processor (such as an IBM 4341), equipped with its own tape units. The total hardware cost of this cluster of emulators and dedicated processor would be in the range of 1.1 to 1.5 Msf, depending mainly on the number of tape units. The potential production capacity of this bank of emulators would be roughly the same as the capacity of an IBM 3081K mainframe, which would cost some 9 Msf with equivalent peripherals. The duty factor might be in the 70-80% range, and the extra tape mounting load would be some 1000 tapes per week.

Although some assistance with development might be available from CERN we believe that banks of 3081/E emulators should only be made available to collaborations which are willing to invest heavily themselves. In particular the collaborations should pay for and operate the dedicated processor, including all tape mounting. They should be responsible for all maintenance of the 3081/E hardware and of the connection to the dedicated processor. Many of the remarks concerning "private" computer centres made in Section 7.3.1 also apply here.

7.3.4. RIOS replacement

There are 10 Remote Input Output Stations (RIOS) strategically located across the CERN site. They can be used to print output generated either at the computer centre or on any of the computers connected to CERNET. However they are now eleven years old, maintenance is expensive, and spare parts for the printers, which have now been out of production for three years, are starting to pose a problem. At the end of 1982 3 RIOS were suppressed and any further suppressions would cause significant problems for remote users. A study is in progress to find a suitable replacement.

> Recommendation 12
> We recommend that the existing Remote Input Output Stations (RIOS) should be replaced as soon as possible by modern low-cost print stations.

The cost would be about 1.0 Msf and 1 manyear of staff time would be needed. It might be possible to spread the costs over two to three years.

7.3.5. Peripherals for the Central Computers

The Working Group 3 report (Appendix D, Sections 5.3.4 and 9.1.2) discusses the key role played by the mass storage system (MSS) in supporting all of the computer centre mainframes, and also very many of the computers attached to CERNET. The report also points to the under-exploitation of central processor resources that is being caused by the absence of an adequate set of peripherals. Experience has shown that regular optimisation of the peripherals attached to a mainframe can be an economic and effective way of improving the overall performance.

> Recommendation 13
> We recommend that, in order to allow for the efficient exploitation of the presently installed processors, the mass storage system should be upgraded, and extra peripherals, including magnetic tapes and disks, should be acquired for the computer centre systems.

This would require the expenditure of 1.3 Msf in 1983, and 0.7 Msf in 1984 and subsequent years.

7.3.6. File Server

As pointed out above, the mass storage system (MSS) plays an important role as a file server for the majority of computers of all sizes on CERN site. Many of the files passing through the central "IBM" computers to and from the MSS are logically independent of the "IBM". (We write "IBM" to indicate the computer centre mainframes supplied by either IBM or IBM plug-compatible suppliers). However the cost of this traffic in "IBM" processor capacity is significant. The throughput is limited by the priority that can be assigned to the transfer task, which is only one of many in the "IBM". We believe it is necessary to dedicate a specific processor to the file server task, both to avoid waste of processor power and to make a necessary improvement in the throughput. A detailed technical study is needed to establish the size of this dedicated machine.

> Recommendation 14
> We recommend that the present mass storage system should be converted into a generalised file server.

The cost of this conversion would be 1.7 Msf spread over the years 1984-1986. The equivalent of one manyear per year of systems programming effort would be required.

7.3.7. Microprocessor Support

We agree with the recommendations of the Microprocessor Working Groups (see Appendix E) that the Motorola 6809 and 68000 should be CERN's preferred microprocessors and that G64 and VME should be the preferred bus systems. We agree with the proposals to provide a central microprocessor support service.

The Microprocessor Report (Appendix E) gives detailed estimates of the money and manpower required for hardware and software support. The money needed is 1.6 Msf in 1983 and 600 Ksf per year subsequently. In total about 20 staff will be involved for a major part of their time in the Microprocessor Support Project. About 14 manyears per year are committed to this sort of activity at present. At least 4 manyears per year of effort must be found to provide general hardware support for the 16-bit systems.

7.3.8. Computing for Engineers

In general the level of computer support for engineers at CERN is inadequate and needs to be improved. This is already being tackled in the fields of Microprocessor Support and Computer Aided Mechanical Engineering. The ad-hoc working group to study Electronic Computer Aided Design and Manufacture has now produced a report. It is clear that with reasonable investment the productivity of the present electronics staff can be significantly improved, and that this will be especially important in the areas of LEP machine construction, Fastbus systems, and trigger processors and emulators.

> Recommendation 15
> We recommend that CERN should purchase a modern computer system for electronics design as soon as possible. This system should be capable of supporting design tools for VLSI (Very Large Scale Integration) chip design and manufacture.

The cost for Electronics Design pilot project will be 300 Ksf in 1983. Full implementation might require a capital expenditure of up to 3MSf spread over 3 years. We further believe that more general action to acquire better software packages for engineers would need an extra 500 Ksf per year. Up to 3 people could be involved in providing a Program Enquiry Office type of service for engineers, including the definition and maintenance of appropriate libraries of software, but these people would have to be found by redeployment of existing staff. Neither the extra 500 Ksf nor the 3 people have been included in the resource requirements that we discuss in Chapter 8.

7.4. Interactive Computing

7.4.1. General Discussion

We remark as a preamble that interactive computing is one of the most important, technically complex, and contentious topics that we have had to deal with. There is a steady development of new products, including new software and even new concepts by the manufacturers, and new user needs and problems with the existing services come to light almost overnight.

Much particle physics computing needs major amounts of processor power allied to significant input/output performance. An important fraction of the load is of a repetitive nature, calling for the similar analysis of many events. This sort of load has always been handled best by computers running batch operating systems, and the CERN computer centre was built on computers supplied by manufacturers who are pre-eminent in that market. This load will continue to grow in the future.

Besides this production line computing, people need to be able to interact with running programs, to change parameters, make tests and see the results without waiting for long printouts to be produced. Some modern operating systems provide good interactive facilities, and can run effectively both on computers whose primary purpose is to provide a batch processing service and on smaller machines devoted to interaction. Experience shows that for many types of interaction the separate machine is more flexible and efficient than the addition of interactive facilities to heavily loaded batch systems. However, even interactive access to those central computers that primarily provide batch processing services is increasing in importance, both for control of the batch production and for access to central data storage.

7.4.2. Possible Approaches to Future Interactive Systems

At least three different solutions have been proposed for the provision of interactive computing services at CERN over the next five years. The first calls for an immediate strong commitment to "personal" workstations, the second calls for the installation of VAX computers running the VMS operating system to provide a general service, and the third calls for the installation of the VM/CMS operating system on some part of the central "IBM" computers.

These solutions all have various advantages and disadvantages, and the field is evolving extremely quickly. None of the solutions, on their own, offers a satisfactory solution to **all** of CERN's requirements. More details are given in the following sections (7.4.3 to 7.4.5).

> Recommendation 16
> It is now too early to decide on a long term strategy for interactive computing at CERN. We recommend that CERN remains flexible and makes trials of various solutions, so that it can continue to meet the needs of the physics programme as the market and the user needs evolve.

The following three sections discuss initial investment in different approaches to interactive computing. By 1985 CERN should have a much clearer idea than now about how it should share its investments between different approaches. Whatever solutions are adopted the total investment required is likely to be significant.

7.4.3. "Personal" Workstations

"Personal" workstations are designed to be used by only one person at a time, thus avoiding many of the problems associated with resource sharing that occur when running interactive systems on larger computers. They typically offer significant computing power, an excellent graphics display and associated input devices, and a sophisticated operating system. Some of the present models already have good connections over the manufacturer's network to file servers and other common services. However, because of a number of factors to do with economics and support manpower, they cannot possibly meet all of CERN's needs for interaction today, or even in the medium term.

> Recommendation 17
> We recommend that CERN should not make any immediate heavy investment in "personal" workstations, but that the present evaluation project in the Data Handling Division should be supported, since we believe that "personal" work-stations are the best medium term solution for the provision of many interactive computing services.

Investment of 200 Ksf per year during 1983 and 1984 will be needed. The present staff effort of about 2 manyears per year should be maintained for the evaluation. In later years it is likely that the expenditure will rise rapidly.

7.4.4. Program Development on VAX Computers

Several VAX computers have already been installed as part of the computer centre for specialised services, such as database or computer aided design work. The VMS operating system has a very attractive interface for the physicist. Many VAXes have been installed in European institutes, and many physicists have asked for a VMS service to be made available at CERN. On the other hand the processor power of a VAX is rather limited and it would not be sensible to base any large (100 simultaneous user) central interactive system on the present models.

> Recommendation 18
> In order to provide an interactive service for those who need it most urgently, we recommend that CERN should provide a program development service for the LEP collaborations based on VAX computers in the computer centre, running the VMS operating system. Resource allocation should be managed by the LEP collaborations themselves.

This will start to offload some major users from WYLBUR, and allow us to gain experience of running VAXes in a more general environment. It is important to understand that we are not proposing to run a general purpose service on these machines, which would imply giving full support for a third manufacturer's equipment beside IBM and CDC. The startup cost of such a service for the LEP collaborations would be 2 Msf based on 2 VAX-11/780s, spread over 1983-1985. Central support for such a service would require about 1 manyear per year.

7.4.5. Interactive Computing on the "IBM"

When CERN's first "IBM" mainframe, the 370/168, was installed in 1976, IBM offered no interactive system that was both adapted to CERN's needs and compatible with the MVS batch operating system. As a result CERN decided to use the WYLBUR semi-interactive system, which had been developed at SLAC. Today WYLBUR, which has been considerably developed at CERN, runs on the Siemens-Fujitsu 7880 mainframe and batch jobs submitted by WYLBUR users are run on the IBM 3081K mainframe. During the daytime these computers are essentially devoted to running WYLBUR and the flood of short batch jobs generated by the users trying to simulate an interactive service. The present WYLBUR service supports 180 simultaneous users, and 1500 people logon at some stage during each week, with 1000 people submitting at least one batch job. To give an idea of the cost of this service, the capital value of the share of the "IBM" resources devoted to this service is about 4 Msf.

It is clear that WYLBUR is by now stretched to its technical limits, and to extend it further would be very costly in manpower and would only exaggerate its obsolescence and incompatibility with forthcoming IBM operating systems versions to which it must be interfaced. However the number of simultaneous users of WYLBUR has grown linearly since its introduction, has doubled in the last three years, and the growth shows no sign of slackening off. We believe that, even if "personal" workstations and specialised interactive services are introduced on a significant scale the growth in demand for access to the general "IBM" services will remain strong until at least 1986. CERN will therefore be faced with a serious problem of providing enough access to these central services during the LEP construction period.

The conclusion of the Working Group, after much study, is that the best step is to move from WYLBUR to the standard IBM operating system VM/CMS. This has genuine interactive facilities, which are more efficient than WYLBUR Exec files, and which, after an initial conversion period, will be more easily adaptable to the CERN users' needs. In the longer term there will be major advantages in moving to a manufacturer supported system. The initial conversion is not, however, trivial, as SLAC and Rutherford have both found out, but CERN should be able to benefit from their work. The first step should be to try out

VM/CMS with as much help as we can get from outside institutes, before planning any major installation.

> Recommendation 19
> We recommend that the Data Handling Division should undertake during this year a full technical evaluation of the VM/CMS operating system in order to be able to make detailed proposals as to how it should be used at CERN.

Since any serious evaluation will take six months, and since it should not take place in competition with the already overloaded WYLBUR service, CERN will need to acquire a small IBM compatible computer for this work. In the medium term this computer, in upgraded form, could be used as the file server computer (see Section 7.3.5). The capital expenditure will be 400 Ksf, which could be taken from the file server budget. The machine should be acquired as soon as possible. During the evaluation the equivalent of 5 staff (systems programmers, operations and user support personnel) will be needed for 6 months.

7.5. Questions of Organisation

7.5.1. Interdivisional Cooperation

It is clear to us that the use of computing and related technologies continues to spread throughout all of the CERN divisions. Efficient use of resources requires that the divisions should try to cooperate in this area. Some of this cooperation can be arranged through projects which are centred in the Data Handling Division but which have good contacts with the users and with staff working in the same field in the other divisions. However this structure is sometimes not appropriate and the job of ensuring that there is sufficient cooperation devolves onto the Director responsible for computing.

It is our opinion that, because of this increasing use of computing in all divisions, there is a serious risk that the Director responsible for computing, who is also in charge of major parts of the physics research program in the present CERN structure, cannot devote enough time to his computing responsibilities.

> Recommendation 20
> We recommend that the Director responsible for Computing should appoint a Deputy for Computing and Communications. We envisage this person being responsible for ensuring inter-divisional cooperation on a range of computing topics, including, for example, office automation, various special purpose computing projects, and networks. In addition he would be the natural person to organise the inter-laboratory collaboration required for a successful European particle physics network.

7.5.2. Collaboration with Outside Institutes

Several outside institutes, notably but not exclusively Rutherford and Saclay, have teams that have general support functions for electronics and/or computing. It has been suggested that some of these groups might wish to take on the responsibility for the coordinated development and subsequent general support of certain common elements of the LEP program. Parts of the online data acquisition software, and consistent support outside CERN of a common particle physics program library were the topics most frequently considered, although collaboration on general electronics has also been mentioned.

> Recommendation 21
> We recommend that the CERN Directorate should approach the managements of the institutes concerned to discover whether they would be prepared to become responsible for the coordinated development and subsequent general support of certain common elements of the LEP computing and electronics program.

7.5.3. Data Acquisition - Detailed Planning

We note that the data acquisition systems at LEP will involve a very widespread collaboration between staff at CERN in many groups (DD/Online, DD/Emulators and Trigger Processors, DD/Network, EP/Electronics, EP/Instrumentation, various EF groups) and members of the LEP groups themselves. E. Gabathuler has already arranged for one member of the DD/Online group and one member of the EP/Electronics group to be attached to each LEP collaboration.

> Recommendation 22
> We recommend that close contact should be maintained at all times among the CERN support teams and the LEP groups in order to continue to identify areas of common interest and to foster joint developments.
>
> Implementation of the recommendations on Data Acquisition, including the definition and execution of the detailed work plans, would, under the existing CERN structure, be the joint responsibility of the Heads of the DD/Online and EP/Electronics Groups. We recognise that this division of responsibility is not an ideal situation, and we suggest that it should be carefully reviewed in the longer term.

7.5.4. Offline Data Processing Facilities

In the LEP context, support for offline data processing will be most meaningful in the area of well defined support packages, of which the GEM memory manager is an example, and in the area of well defined support facilities, of which the Patchy source management system is an example. Other areas which are growing in importance for physicists, and where a concentrated support effort might be needed in order to provide the required facilities, include advanced interactive graphics, data base management, and user transparent network interfaces.

Owing to the size of the LEP collaborations, standard practices will have to be adopted in many areas of the data processing. These standard practices will either involve the use of accepted tools, or the adoption of a set of clearly stated rules. It will be very important that data processing experts from the DD and EP divisions participate in LEP experiments and that they foster effective collaboration between the experiments and also among the whole European particle physics community.

8. RESOURCE REQUIREMENTS

8.1. General

The fact that this is the first report for a considerable time to consider computing and electronics needs across all CERN divisions has meant that there are some areas where it has been difficult to assemble figures for the past and present staff and money commitments. Nevertheless we believe that we have now reached a good understanding of the real situation in almost all areas.

Each of the recommendations presented in Chapter 7 has been followed, when appropriate, by a discussion of the staff and money that will be required for their implementation. We would like to emphasise that we all understand the present CERN budget constraints, and that our estimates have been based on a **minimum** investment consistent with being able to take and analyse data successfully at LEP.

The question of which Divisional budget should fund any given recommendation has not been addressed. According to previous CERN practice the Data Handling Division would be responsible for most of the General Facilities and Communications, but for very little of the expenditure on Data Acquisition and Electronics, which would be covered by EP and the Accelerator Divisions.

In a similar way we have given the total number of extra staff needed to implement each recommendation, but have not discussed where these staff might come from. There is certainly no requirement for all of these people to be CERN staff members. Providing that they have the right background to fit the job they could come from a number of sources, including Fellows and Research Associates (these are short-term CERN appointments), external institutes, or the LEP collaborations themselves. Nevertheless it is important to realise that coordinating the activity of a number of short-term staff can itself be a very time consuming occupation, and that projects do need a core of long-term expertise.

We have assumed that LEP experiments will themselves pay for all hardware components which are supported by CERN teams but which they need in quantity. For example this applies today to CERNET links, terminals, trigger processors and Camac modules. Fastbus modules, Local Area Networks and "personal" workstations will almost certainly come into this category in future.

There are some areas of activity, such as trigger processing and emulators, which are important but for which we give no staff and money estimates since their development and support are adequately staffed and funded at present.

Compared to some of the figures given in the Working Group reports we have taken advantage of the 6-12 months delay in LEP completion (with respect to the end-1987 date), recently announced by the Director General, in order to lengthen the period for full implementation and payment of our recommendations to include 1988. There are related beneficial effects in slightly reducing some of the staff numbers.

In the following sections we discuss the recent (1980-1982) level of CERN funding, the present CERN budget levels, and the resources that will be needed to implement our recommendations in the three areas of General Facilities, Data Acquisition and Electronics, and Communications. All discussion is in terms of 1983 prices.

8.2. General Facilities

Averaged over 1980-1982 CERN has spent 15.1 Msf per year on the Computer Centre and general interactive systems, of which 6 Msf per year was capital expenditure. The present funding foreseen for 1983-1988 is at the level of 9.0 Msf per year, with no allowance for capital expenditure. In view of the uncertainty related to the installation of "private" computer centres, and related to the total central batch capacity that will be required for LEP, we have not included any estimate for the money that CERN needs to budget in this area. Our recommendations necessitate the following capital expenditure

1983	1984	1985	1986	1987	1988
3.1	3.1	2.7+	2.0+	1.7+	1.7+

The + sign for 1985-1988 is used to indicate the significant, but at present undetermined, sum of money that should be spent on interactive systems (see discussion in Section 7.4). The precise implementations will depend on the experience obtained with the alternative approaches, as discussed in Section 7.4.

The present staffing of about 66 should rise by about 4 in order to cope with the increased diversity, complexity, and decentralisation of the systems that are being supported. Unless magnetic tape handling in the computer centre can be held at the present level some increase in the number of service contract staff will be needed, corresponding to an additional sum of 500 Ksf per year by 1986.

8.3. Communications

Averaged over 1980-1982 DD Division has spent about 1.0 Msf per year on CERNET, and both SPS and PS Divisions have also spent money on communications for their control systems. Roughly 50% of the total has been capital expenditure. The requirements of the LEP project for communications related to the control system are assumed to be adequately covered in the existing plans. The present funding foreseen for CERNET during the period 1983-1988 is at the level of 500 Ksf per year for operation, with no allowance for capital expenditure.

Our recommendations would require the following expenditure

1983	1984	1985	1986	1987	1988
0.2	0.7	1.1	0.9	0.8	0.6

The present staffing of about 13 needs to rise by 10 to a peak of 23 in 1985. The peak is caused by the requirement to keep the present CERNET facilities operating smoothly while introducing and integrating the new backbone network, gateway connections, Local Area Networks, and external connections.

8.4. Data Acquisition and Electronics

In view of the wide involvement of outside institutes in these activities at CERN we have not attempted to give figures for the total money and manpower used in this area. Our recommendations would require allocation of the following sums for capital expenditure on projects which either have no foreseen budget in the present CERN plans, or where the present budget is clearly inadequate. In the second case the numbers are the extra expenditure that is required.

	1983	1984	1985	1986	1987	1988
Data acquisition support	0.1	0.1	0.1	0.1	0.1	0.1
Test systems (Caviar-like)	0.4	0.4	0.4	0.4	0.4	0.4
Fastbus support	0.2	0.4	0.4	0.4	0.2	0.2
Microprocessor support	1.6	0.6	0.6	0.6	0.6	0.6
Electronic CAD	0.3	0.8	0.8	0.6	0.6	0.4
Total	2.6	2.3	2.3	2.1	1.9	1.9

Compared to present staffing levels 14 engineers and technicians will be needed, 4 for Fastbus development, 4 for Fastbus support, 4 for microprocessor support, and 2 for the Online Computer Pool. In addition 6 extra programmers will be needed, 2 for Fastbus development and 4 for data acquisition systems support.

8.5. Conclusions

Electronics and computing facilities are just as important a part of today's physics experiments as are the accelerators and detectors. CERN has recently demonstrated unequivocally the benefits, in terms of physics results, that have followed from being able to properly acquire, distribute and analyse the data coming from the $p\bar{p}$ collider at start-up. We have presented a set of recommendations for computing and electronics projects that we consider will be necessary if LEP data is to be exploited in a similar way. We note that both because of the high requirement for Monte Carlo studies and because of the unknown quality of the background when the LEP machine first runs, the effect of any initial low luminosity will not significantly reduce the need for raw computing power.

Leaving aside the question of the increase in the central batch capacity, for which technical and political decisions are required before a valid budget model can be prepared, there is a need to spend about 6 Msf per year between now and 1988, and some 30 additional posts will be required. These figures may seem rather high in the present CERN budgetary context, and we would like to make several explanatory points:-

- We have reviewed the estimates given for each recommendation, in order to ensure that they are as low as reasonably possible.

- It must be realised that the totals cover several projects which are of direct benefit to the accelerator divisions.
- The scale of the expenditure should be seen in relation to the expenditure of 900 Msf on the construction of LEP and 300 Msf for the four LEP experiments, and to the major benefits that will result in terms of physics productivity.
- Whereas the costs of equipment are continuing to fall, the manpower needed, particularly for communications and for direct support of physics groups with advice, maintenance, etc., is not falling in the same way. Our recommendations are aimed at minimising the manpower needs, but CERN will not escape the world-wide trend of having to increase the fraction of its staff involved in computing and communications activities.

We fully recognise that there may be a real problem to match the resources required to those available at CERN. In order to carry out this programme it will be necessary to find the manpower mainly at CERN, with help coming from outside institutes and the experiments themselves. The possibility of providing Research Associate and Fellowship positions for well-defined project work should be expanded. It will also be necessary for the CERN Division Leaders, in particular the leaders of the Research Divisions, in collaboration with the Director responsible for Computing, to review their present development and support activities in order to find the staff required for implementing the proposed programme. In many areas the recruitment of young staff with backgrounds in modern electronics and computing will be important if we are to maintain a successful CERN physics programme at the end of this decade.

9. REFERENCES

The background to each of our recommendations can be found in the Working Group reports, as indicated below.

Our recommendation	Working Group discussion	
1	WG1	Conclusion 2; p.15/17
	WG2	1.1; 1.2.6/11; 2.1; 3.1/4
2	WG2	2.2.4; 3.5; Appendix 4: 1; Appendix 5
3	WG2	1.2.2; 1.2.12; 3.6; Appendix 6: 7
4	WG1	Conclusion 2; p.17
	WG2	1.1.2; 1.2.4; 2.2.1; 3.1; 3.7
5	WG2	1.2.6/9; 2.2.3; 2.2.6; 3.3; Appendix 4: 2.5; Appendix 7: 6/9
6	WG1	Conclusions 2,8; p.15/16
	WG2	1.2.6; 2.2.2; 3.2; Appendix 4: 2.2/5
7	WG2	1.2.5; 2.2.5; Appendix 5: 3.2/3
8	WG1	Conclusions 7,12; p.13/14
9	WG1	p.7; p.17
10	WG3	9.1.10
11	WG3	Recommendations 2,3; 3.4; 9.1.1; 9.1.10
12	WG3	9.1.5
13	WG3	Recommendations 11,12; 5.3.4; 9.1.2; 9.1.4
14	WG3	Recommendation 13; 5.3.4; 9.1.3
15	WG1	Conclusion 11; p.19
	WG3	8.3
16	WG3	Recommendations 4/8; 5.4; 6
17	WG1	Conclusion 6; p.8
	WG3	Recommendation 5; 6.4.3; 9.1.9
18	WG3	Recommendation 6; 6.4.2; 6.6
19	WG3	Recommendation 7; 6.6/7; 9.1.6
21	WG1	p.18
	WG2	3.1
	WG3	Recommendations 9,10; 5.4
22	WG1	Recommendation 1; p.18

10. GLOSSARY

In this glossary we try to list the full meaning of all abbreviations used in the LEP Computer Planning Group Report. We also try to explain some computing jargon for physicists and some physics jargon for computer specialists.

CAD: Computer Aided Design. The use of specially developed hardware and software to help in the engineering design process.

Camac: A manufacturer-independent bus system widely used to interface particle physics experiments to computers.

Caviar: A CERN designed system, based on the Motorola 6800, and used for testing electronics and detectors, etc.

CCITT: Comite Consultatif International pour les Telephones et Telegraphes. A committee of the International Telecommunciations Union which sets the international standards for PTTs and other suppliers of communications systems.

CDC: Control Data Corporation. A computer manufacturer.

CERNET: An onsite high-speed packet switched network developed by CERN during the second half of the 1970s. Connects about 100 computers from many different manufacturers. Paths of up to 10 km. are already provided, and this will have to increase when LEP is built.

Collaboration: The people who collaborate to carry out an experiment. Large experiments may have 10-20 institutes collaborating, with a total of 200-300 staff. The experiment might take 10 or more years from initial planning to final data analysis.

DD: Data Handling Division at CERN. Although this is the main "computing" division, there are several strong "computing" groups in other divisions.

DEC: Digital Equipment Corporation. A computer manufacturer.

Directorate: The top level management body at CERN. At present consists of 1 Director General and 5 Directors.

Division: The normal administrative unit at CERN, typically with 200 to 500 staff.

ECFA: European Committee for Future Accelerators. This body, whose principal role is to take care of Europe's requirements for future particle accelerators, has also looked at particle physics data handling on a European-wide basis.

EF: Experimental Physics Facilities Division at CERN. Responsible for many of the large "facilities" used for physics.

Electronics Pool: A service in EP Division for the selection, testing, repair and loan to physics groups of a wide range of electronics equipment.

Emulator: A microprogrammed device that emulates the machine instructions obeyed by a larger computer.

EP: Experimental Physics Division at CERN.

EXEC Files: Files containing sequences of WYLBUR commands.

FASTBUS: A specification for a very high speed bus developed for use with data acquisition electronics. In many respects a natural successor to Camac.

Fellow: A short term (typically 2-3 years) CERN appointment. Fellows are normally younger than Research Associates.

File Server: A computer dedicated to maintaining a file base on behalf of other computers accessible over a network.

GEM: A memory and data management package for FORTRAN programs.

G64: A bus system (principally) for 8-bit microprocessor systems.

G700: A set of CCITT specifications for the transmission and multiplexing of digital signals.

HP: Hewlett Packard. A computer manufacturer.

IBM or "IBM": International Business Machines. A computer manufacturer. We write "IBM" when we wish to indicate computers which are either made by IBM or which are exactly compatible with IBM computers, but nevertheless made by other companies.

IEEE: Institute of Electrical and Electronics Engineers. An American institute which has taken the responsibility for many bus standardisation efforts.

Index: A circuit switched data communication system extensively used at CERN to connect terminals to computers. The terminal user can dynamically select the computer to which he wishes to be connected.

Local Area Network (LAN). A communications network which is geographically limited (typically to a 1 km. radius). These networks are usually of relatively high speed and should allow for easy interconnection of terminals, microprocessors and computers within adjacent buildings.

LEP: Large Electron Positron Collider. A 27km circumference accelerator being built at CERN, which will bring bunches of electrons and positrons into collision.

Mainframe: A computer that is primarily used for processing data, rather than for acquiring data.

Minicomputer: Historically, cheaper computers used for dedicated functions. The first minicomputers normally had 16-bit architectures, but there are now many 32-bit minicomputers available. Many of the dedicated functions are now carried out by microprocessors, and the features of the 32-bit minicomputers now rival those of the previous generation of mainframes.

Microprocessor: Processors which are very cheap and physically very small. Can be integrated with terminals and storage to provide microcomputers, which are then very cheap computers, or integrated into special purpose electronic equipment for control etc.

Monte Carlo: Calculations carried out using the techniques of random number generation. In particle physics these can be useful in order to understand the behaviour of the experimental apparatus, or to understand the (statistical) physical processes being studied.

Motorola: A manufacturer of semiconductor chips and microprocessors.

Msf: Mega Swiss francs.

MSS: Mass Storage System. The IBM 3850 cartridge tape system.

Online Computer Pool: A service in DD Division which deals with the loan and installation of online (at experiment) computers, with their integration into CERN networks such as Index and CERNET, and which tries to ensure that the whole set-up remains operational for data taking.

Packet switched: Describes a network in which the information flowing from a source to a destination is carried as a set of "packets". Blocks of information are split into separate packets at the source and reassembled at the destination.

PATCHY: A source code maintenance program developed at CERN.

$p\bar{p}$ collider: A 7 km. circumference accelerator at CERN which collides bunches of protons against bunches of antiprotons. Where the first signs of the intermediate vector boson (W) have recently been detected.

Print Server: A computer dedicated to the function of printing files on behalf of other computers.

Production: Large scale analysis of the data from a physics experiment. This "production" processing may require the mounting of thousands of magnetic tapes, and can occupy even the most powerful computers for weeks on end.

Protocol: An agreement about how to transmit data, especially across networks. Low level protocols define the electrical and physical standards to be observed, and deal with the transmission and error detection and correction of the bit stream. High level protocols deal with the data formatting, including the form of messages, the terminal to computer dialogue, files, etc.

PS: Proton Synchrotron. A CERN accelerator and also a CERN division.

PTT: Postes, Telephones et Telegraphes. The people who run the national telecommunications services.

Personal Work Station. A computer sufficiently inexpensive to be dedicated to a single user. We distinguish a Personal Work Station from a Personal Computer (PC) by its greater power.

Research Associate: A short term (typically 12-18 months) CERN appointment. Normally older than Fellows.

RIOS: Remote Input Output Station. Equipment for reading cards and printing output at remote locations on the CERN site. The card reading role is largely historic.

Rutherford: The Rutherford Appleton Laboratory, a major British laboratory involved in particle physics research.

Saclay: A major French laboratory involved in particle physics research.

SLAC: Stanford Linear Accelerator Laboratory. A major U.S. laboratory involved in particle physics research.

SPS: Super Proton Synchrotron. A CERN accelerator and also a CERN division.

Trigger processor: A processor whose task is to decide whether the data from the initial "trigger" of a particle detector are of any further interest or not. Since selection by a factor of 10,000 is not uncommon between the initial trigger signal and data recording, many different trigger processing techniques are used, often in several levels.

VAX: The name given to a range of 32-bit computers manufactured by DEC.

VLSI: Very Large Scale Integration. Refers to semiconductor chips composed of very many tightly packed logic elements or memories.

VM/CMS: Virtual Machine / Conversational Monitor System. An IBM operating system running on 43xx and 30xx series machines, providing efficient support of large numbers of interactive users.

VME: Versabus Modules Europe. A specification for attaching equipment to microprocessors, especially the Motorola 68000 family. Uses Euromechanics and a 16/32-bit bus definition. Supported by Mostek, Philips/Signetics and Thomson-EFCIS as well as Motorola.

VMS: The virtual memory operating system for DEC's VAX computers.

WYLBUR: A conversational editor, file management and job entry system running under the MVS operating system on IBM computers.

X25: A CCITT recommendation for packet switched communication over public networks.

Appendix A

LCPG MEMBERSHIP

M.C.Crowley-Milling/LEP	
R.Dobinson/EP	Chairman WG1
E.Gabathuler/DI	Chairman LCPG
J.M.Gerard/DD	Chairman WG2
H.Grote/DD	
M.G.N.Hine/DD	Secretary LCPG
P.G.Innocenti/SPS	
C.M.Jones/DD	Chairman WG3
A.Wetherell/EP	
D.Williams/DD	
B.Zacharov/DD	
P.Zanella/DD	

CONTENTS

APPENDIX B
TO THE LEP COMPUTING PLANNING GROUP REPORT

THE LEP COMPUTING PLANNING GROUP ON DATA ACQUISITION AND MONITORING.

Final Report

15th January 1983

Introduction

The purpose of this working group has been to investigate the detailed future requirements for on-line data acquisition and monitoring with special emphasis on LEP.

It should be stressed that this report is of necessity preliminary in nature. It is merely phase 1 of a multiphase process. Many discussions are still required, and many topics remain to be explored in greater depth. LEP groups themselves are only just beginning to formulate their detailed plans. Despite this there are a number of apparent needs which seem to have become clear in our discussions, and it is these which are emphasised in this document. In most cases further work of a practical nature is required during 1983 to clarify more precisely outstanding technical points and obtain experience in new fields. We include suggestions for the form such work could take and recommend that CERN's contacts with LEP groups are strengthened and put on a longer term basis in order to ensure efficient and cost effective support in the years ahead.

Conclusions and recommendations for future work

1) Discussions with LEP groups will need to continue throughout the period prior to the start up of the machine and a mechanism to foster such discussions is required. Contact must be established with outside laboratories to encourage collaboration in the field of electronics and computing. No formal mechanisms exist for this at the moment . This would help avoid duplication of effort and increase efficiency and cost effectiveness. Most LEP groups are in a very preliminary planning stage, they may in some areas not yet know what they want, or even know what is available to them. We feel it is therefore too early to formulate computing and electronics needs beyond what is outlined in this report

and suggest that such an aim can only be realised by far more detailed and wide ranging work involving LEP collaborations and CERN service groups, in parallel with practical experience of the new technologies which are likely to be used in future experiments. We recommend the establishment of technical study groups comprising LEP experimenters and CERN support services. These study groups should further investigate in depth requirements for triggering, readout, monitoring and communications. Where possible standards should be agreed, clearly defined and supported by CERN. It should be realised that this work will take many months. Until we understand which requirements are common to several experiments it will not be possible to decide precisely which facilities and resources should be provided centrally and which allocated to individual experiments.

2) Networking and communications are perhaps the most important single computing topic at LEP and we stress the need for an integrated approach to these problems. Extensive support is required for world wide HEP networking. Much of this effort must come from outside CERN and we recommend that CERN together with ECFA and outside laboratories continue to work together towards providing such facilities. Access is required as soon as possible from many more outside institutes to computers on the CERN site. Good CERN wide communications will be essential to most people, physicists, engineers, machine builders, administrators, etc, they must be tackled in a coherent way throughout the organisation. A high speed backbone network is required to allow physicists to communicate over the whole expanded CERN site. At the level of individual experiments we see a need for CERN to support one or more "standard" local area networks. We recommend the acquisition of this type of equipment as soon as possible and that a program of tests be carried out to clarify the optimum approach for LEP experiments. It is important that whatever is chosen as a CERN standard should enjoy widespread commercial support.

3) LEP groups wish to standardise within a collaboration on one family of 32 bit minicomputers. If all experiments were to use the same computer clear economies in the manpower and services provided by CERN would result. CERN should review its policy of supporting both DEC and NORD minicomputers for data acquisition. In particular whether the manpower savings of a single manufacturer policy would outway the commercial and technical advantages of continuing to develop systems based on computers from both manufactures.

4) The conclusions of the CERN Steering Committee for Microprocessor Support are in general supported in particular the proposed use of the Motorola 68 000 and VME bus. However, WG1 warns against too much polarisation in the fast moving field of micro devices and notes that the provision of FORTRAN on 16 bit machines is essential. We note a need to investigate how the requirements for small test systems can best be met.

5) The role in triggering of special purpose high speed processors, such as the 3081E, XOP etc, and large arrays of conventional microprocessors needs further study. We support the development of both approaches for the future and suggest more practical experience in physics experiments over the next year would be very useful as a guide to LEP applications.

6) CERN lacks experience of personal work stations such as PERQ and APOLLO. We think this is a field of relevance and importance for the future and support the purchase of such systems by CERN for evaluation during 1983. We see a particular interest for the physics community in applications such as control and monitoring of experiments (including interactive graphics), program development and computer aided design. However we emphasise that at this point in time the field is expanding rapidly and that firm decisions can probably only be taken at the end of 1983 in the light of accumulated experience.

7) FASTBUS should be supported as a high performance standard. However there is a need to discuss in detail with LEP groups to understand how and where it can best be used i.e. a process of mutual education needs to take place leading to the definition of a range of useful modules. We feel further small FASTBUS projects should be encouraged in the immediate future and stress the importance of CERN providing general purpose devices such as crates, protocol ICs, segment interconnects and computer interfaces. We recommend that CERN actively explore problems of large multi-crate, multi-computer systems and the interconnections of crate segments over distances of up to several hundred metres.

8) Experiments will use a variety of buses and networks (CAMAC, VME, FASTBUS, ETHERNET, etc.). We recognise the need to study and systematise their interconnection.

9) The trend towards distributed computing, although bringing many obvious benefits, can lead to an undesirable proliferation of different systems. We recommend that transportable standard software packages and systems should be made available where ever possible.

10) At the moment it seems likely that data will be recorded at the experiments using 6250 bpi magnetic tapes. However, we suggest developments in recording technology be monitored carefully as new products may become available in time for LEP. The final decision on how and where to record raw data does not need to be taken yet and it is important that future network planning must take account of possible remote recording.

11) The field of electronics is fast moving and ever more complex. WG1 recognises the need to use advanced state of the art techniques in future generations of experiments. Computer aided design (CAD) is a prerequisite for efficient electronics development and production and we recommend that CERN should acquire a modern electronic CAD system as soon as possible. We support the need to actively investigate the possible role of custom designed ICs.

12) Collaboration with industry, and in particular European industry, should be fostered. We feel a closer involvement by industry, for example joint development projects, would allow better and cheaper products to be made available to the physics community.

An overview of the on-line facilities required by LEP experiments

LEP experiments will be enormous, both in terms of the number of physicists involved and in the complexity of the apparatus. The construction of general purpose assemblies of detectors with very large numbers of channels will mean a wide range of problems to be solved in electronics and computing.

Earlier predictions about large event sizes have been confirmed. Numbers around a few hundred kbytes per event are agreed, although some LEP groups quote even higher figures under certain conditions. This, however, is for real events, but no-one seems to be sure what the event length will be for background triggers, since one cannot yet accurately predict the form these will take. However, LEP groups seem to be prepared to accept they will have to record for off-line analysis of the order of 1-5 Hz events. It is worth noting that 5 Hz at 200 kbytes per event is 1 Mbyte/s, this is in excess of the recording rate of current industry standard magnetic tapes. The bulk of this data will come from the central detectors, although it should be noted that the raw data rate produced by FADCs and CCDs is much higher. There is a clear trend to squeeze the maximum of information from a detector channel.

The data flow and interconnection topographics to be used are important as they are the key factor in determining the level of performance we need from computers, buses and local networks. The subject is also rather complicated given the scale of proposed LEP experiments and groups are only now beginning to consider their needs in detail.

Despite this certain common trends can be discerned:

- The use of electronic subsystems for the acquisition of data from different detectors and for monitoring.

- The independence of subsystems required during the commissioning of different detectors.

- The need to interconnect and integrate subsystems, developed in many different institutes, during the final phase of setting up the experiment and for data taking.

- The need for continuity throughout all phases of the experiment, such that work done in the laboratory or in test beams can be carried over smoothly into the final set-up.

It is likely that Local Area Networks (LANs) will be used to interconnect subsystems and to integrate their component parts at medium speeds. For higher speed communications data acquisition "networks" comprising large number of interconnected CAMAC or FASTBUS crates will be used. Fig. 1 shows a preliminary layout of the on-line system of a typical LEP experiment.

Whilst LEP groups are slowly realising the magnitude of their monitoring and calibrations problems, there are so far no detailed ideas on how to tackle them. Present experiments are on the edge of what can be handled using current techniques and some sort of quantum jump will be required. One of the biggest difficulties will be to coordinate and synchronise the activities of the large number of processes and processors involved. Further work is needed here to define and solve such problems and to deal with the information produced by large complex systems with highly distributed intelligence. For example pulse height and time spectra from a typical detector may be accumulated by a large array of microprocessors. It will not be possible to scan all such information "by eye"; it will need to be checked automatically for anomalies and archived on a central data base.

A good deal of thought will need to go into the presentation of the results of calibration and monitoring. Of particular importance will be the ergonomics, or human factor, of man-machine interactions. Personal work stations offer one particularly exciting solution to this problem offering a high degree of responsiveness and high quality graphics capable of presenting large amount of information in an easily understandable form.

In conclusion LEP experiments will require the means to coordinate the activities of many different processors, an effective data base management system and a very friendly user interface.

Standard minicomputer support

The support of standard minicomputer hardware and software by CERN is essential. For the future this support should be based on the use of powerful 32 bit machines. There is a clear realisation that the present generation of 16 bit minis will be inadequate for LEP.

LEP groups recognise the potential convenience and economy of manpower obtainable by standardising internally on a single family of minicomputers. There is a wish to avoid an inhomogeneous collection of manufactures and software environments. WG1 sees clear advantages to CERN if all LEP experiments were to adopt one well supported standard; less manpower would be required for hardware and software development, and for maintenance and servicing. If a single standard minicomputer is to be used for LEP experiments it should encompass a wide range of products, including extensive communications equipment, and enjoy, if possible, world wide vendor support and user familiarity.

A reasonable assumption is that groups will want to use either DEC (VAX) or NORD (N500) families. There have been some attempts to compare current offerings from these two manufacturers and to look at the possible evolution of their products, this work should continue. Some comment and criticisms have been made about the suitability of VAX hardware and sofware for data acquisition; in particular the large overhead of the current CERN system to acquire and monitor events. A careful and thorough audit of this problem is required especially as there is apparently contradictory experience from other laboratories. It should be noted that the acquisition needs of small tests are different from those ultimately required in a large experiment. The former demands high rate acquisition of small events, the latter smaller numbers of much large events.

Interest has been expressed in the newest and cheapest VAX family member the 11/730. We recommend CERN should acquire such a machine and evaluate its suitability for a data acquisition role in the future.

The role of personal computers and work stations

Personal computers are one of the most exciting and fast moving technologies of the moment. It is a field with a bewildering choice of products and one in which new offerings appear on the market place almost daily.

WG1 has considered the possible use of PERQ, APOLLO and similar personal computers for data acquisition and monitoring in physics experiments. At the present time we do not see them replacing traditional minicomputers. The main reason is that an enormous investment has been made in specialised hardware and software for specifically high energy physics applications and there is no wish to repeat this work. A further reason is the unproven hardware and software of many personal computers.

Despite the misgivings of WG1 about personal computers playing a major role in on-line computing for LEP experiments there are clear areas where they will be of considerable interest. Physicists will need to be able to interact with computing systems in a user-friendly way, to control their experiment, to do physics. No one can be unimpressed after seeing the high quality bit map raster graphic displays available on PERQ and APOLLO, puck or mouse like devices do offer a high degree of convenience for users. However, it would appear too early to rush headlong into the personal computer business. There are two reasons for this. Firstly CERN is lacking in practical experience, secondly prices and products are evolving rapidly. We very much support the acquisition by CERN of a limited number of personal computers for evaluation, we note however that it is not just the top end of the market (PERQ and APOLLO) which could be of interest and relevance, but also the low end (DEC, IBM, and why not APPLE?). Operating these computers in a practical environment, in particular that of a physics group or application, will allow us to better assess their possible role in the future. Several LEP groups have expressed an interest to achieve similar facilities to those offered by personal computers using the newly introduced VAX station.

Working group 3 has recognised the importance of personal computers in its survey of General (off-line) Computer services, in particular for implementing interactive computing needs such as program development, graphics and computer aided design.

Standard microprocessor support

The use of microprocessors in LEP experiments will be widespread. There will be hundreds, and perhaps even thousands of these devices used for data acquisition and reduction, triggering and monitoring. There is a need to provide in a standard way a wide range of tools and facilities; cross assemblers and compilers, interactive languages, libraries, linkers, development systems etc. In general manufacturers of microcomputer products offer less help for end users than do minicomputer manufacturers. The incentives for standardisation are therefore, if anything, greater.

The activities of the CERN working groups for Microprocessor Standardisation have been followed by WG1 and comments have been given to the steering committee. The main points raised were as follows.

- The CERN initiatives in this area were, in general, appreciated.

- There is concern from outside groups and labs that their special needs and problems should not be forgotten. This is also stressed by the ECFA microprocessor working group who have submitted a detailed criticism of the CERN report.

- Some people in WG1 were concerned that too much weight was being given to accelerator division requirements.

- There is considerable support for the choice of the M68000 as a standard. However WG1 stresses the fluid nature of microprocessor technology and urges against too rigid an approach to standardisation.

- There was no interest in the E3S system. It was felt, if CERN were to adopt a bus standard, then it should be a currently existing European, and preferably world wide, system. For this reason VME had very strong support.

- It was felt that the role of FORTRAN had been insufficiently stressed. A compiler must be provided for any 16 bit CERN standard micro. This could allow the migration of tasks from minis down to micros in an easy way in physics experiments.

- Doubts were raised about the use of PILS as an interactive language on smaller micros. If not PILS then what other language could be used?

- Members of WG1 who came from outside laboratories raised the question of payment for standard CERN supported software. This is a potential problem if packages are purchased commercially by CERN. Outside institutes would like, in this case, some form of global high energy physics licencing.

WG1 notes the current widespread use of the CAVIAR system especially in test and development situations. A similar device with enhanced performance is probably required for the future. Effort is required to specify further this need and the implementation details of such a system; M68000, packaging, interfaces to CAMAC and FASTBUS, connection to local area networks, PILS, etc.

Raw data recording

At the present time remote recording of bulk event data is not planned by any experiment, ie) for LEP data will be recorded at the intersection region. However, it is recognised that there are many advantages in not having to deal with tape mounting in the counting room. In the absence of anything better industry standard 6250 bpi magnetic tape will be the medium used but the door should be left open for new products such as video discs. It should be stressed however that this type of device may be only really of widespread interest and impact if the recording medium is truely transportable, as are present day magnetic tapes. A final remark about raw data recording is that if recording via some new device or medium only became economically viable when located centrally, then there would be an immediate requirement for very high speed long distance links around CERN site. The final decision about how and where to record raw data could be left open for one or two more years.

It is likely that test set-ups for LEP experiments will be interested in recording raw data remotely either at the IBM centre or on computers belonging to the collabo- rations.

Languages

The argument about programming languages continues to rage, FORTRAN is as popular as ever amongst the physics community and there is a strong wish to guarantee the availability of this language across both mini and micro computers. Interpretive languages are very useful for testing and debugging and we note the widespread popularity and use of CATY and BAMBI within the high energy physics community. One of the problems that arises at the moment is a lack of transportability of this type of program to the run time environment of experiments; applications often need to be reprogrammed and this is time consuming and uneconomical. The use of the PILS language is therefore of considerable interest for the future, it offers not only the best of existing interpreters but potentially the ability to compile modules and task build them with object code from FORTRAN, assembler and other sources. PILS should be supported by CERN on standard minicomputers and microprocessors, further effort is required to develop the language, in particular to compile it, and make it widely available.

Other high level languages, in particular PASCAL, may find a use in high energy physics applications in particular for use on microprocessors if FORTRAN is not available, however, further studies will be required before the attitude and needs of physicists can become completely clear.

Transportable software packages for on-line systems

We have not had extensive discussions on this topic, although the strong statement can be made that physicists are going to be very unhappy if they have to learn a large number of different systems. This is going to be one of the drawbacks of highly decentralised computing unless we are very careful. Transportability is therefore vital in particular between on-line and off-line environments. One can summarise the obvious advantages at three levels:

a) Users see the same command structures, terminal interaction and representation of data.

b) Programmers see the same languages and library subroutine calls.

c) The duplication of effort by support staff is minimised.

The use of a universal operating system across the whole spectrum of micros, minis and main frames is an immensely attractive idea. For this reason UNIX is of interest for the future, however, its non-friendly interface, at least for physicists, and its apparent lack of real time facilities argue against its truly universal use.

Special purpose processors

It is being assumed that LEP triggering and data compression will take place on multiple levels and that programmable processors of various kinds will be used for this as well as for monitoring functions. Although precise needs are not yet clear since few groups if any have worked out their requirements in detail, a consensus seems to be emerging for a pipelined tree structure of data acquisition hardware which lends itself well to the efficient use of processors.

For the faster levels of triggering, there is a strong request from one LEP group to study in more detail the implementation of a new range of modular trigger electronics; standard building blocks in hardware (and in software too if programmable steps are included).

We have discussed the role of special purpose processors and large arrays of commercially available microprocessors and at present there seems to be a complementary role for both types of device. Two kinds of special purpose devices can be identified; the truly 'special' processors for the faster levels of triggering (e.g. ESOP, XOP,) and the 'emulators' of large minis and mainframes (e.g. 168/E, 370/E, 3081/E, FFP[VAX/E]) which find a place in the more complex levels of triggering, monitoring and off-line analysis. It remains to be seen if massively parallel arrays of high-performance commercial micros will offer an efficient and economic alternative to either kind of special processor. One thing is clear; future experiments will require very large amounts of number-crunching power to trigger and to process data (including complete event reconstruction on-line). It seems very unlikely that this power can be provided commercially.

We recommend further studies and practical experiments to clarify the issues and choices in this area. Useful inputs should come from the conference on the use of specialised processors to be held in Padua, 23-25th March 1983.

On-line graphics

There is a clear requirement to be able to display and manipulate LEP events in an on-line environment. Colour in graphics systems will become ever more important. How best to provide an adequate facility is unclear, in particular, should one use a specialised device such as a MEGATEK or a high performance general purpose personal computer. This is one of the things that should be studied on CERNs PERQs and APOLLOs over the next year.

Modular electronics and bus systems; CAMAC and FASTBUS

CAMAC will not disappear at LEP or any other experiment for some time yet. There may be a need to boost its performance at certain levels, for example, to provide intelligent controllers at a crate and branch level and to interface more efficiently with minicomputers.

Recent technical progress with FASTBUS both in the United States and Europe has been most encouraging. At CERN a single crate system has taken data last summer in an SPS experiment. Joint developments by CERN and the University of Illinois have resulted in the operation of a four crate, three computer, system over distances of 100 m.

Considerable interest in FASTBUS has been shown by two major commercial companies. One of them has announced development of several modules of potential interest for future experiments; a multi hit TDC, a high resolution ADC and an image chamber analyser.

FASTBUS is being actively considered by all four LEP groups. It seems fairly clear that where CAMAC is too slow (for example in central detector electronics) or where sophisticated architectural features are required FASTBUS will take over. CERN should therefore commit itself strongly to providing basic FASTBUS support:

- crates, power supplies and cooling
- interfaces to computers and CAMAC
- general purpose devices such as segment interconnects, display modules, ancilliary logic, extenders, etc.

- interfacing aids; in particular custom chips to implement the FASTBUS protocol
- software for standard mini and micro computers
- user education and help.

Much of this work has, in fact, already been started. In order to help interested physics groups to evaluate and learn about FASTBUS a "starter kit" of basic hardware and software is being organised by CERN. Interfacing to computers and CAMAC is being actively reviewed.

Interest by European industry needs to be strongly encouraged to ensure the commercial availability of equipment.

We recognise three areas of FASTBUS where further technical work is necessary:

1) Close collaboration with US laboratories is required to define and produce protocol chips as soon as possible.
2) A study should be made of high speed long distance interconnection paths between FASTBUS crates. Connections, probably serial and parallel, over distances of several hundred metres will be required at LEP.
3) The results of the review of interfacing computers to FASTBUS are urgently required; should one access FASTBUS via CAMAC or directly via the computer I/O bus? A policy must be defined and a commercial supplier of equipment found. There exists a general question of how to connect buses of different types, and also of how to connect buses to local area networks. The need to try to systematise such intercouplings is clear, but little or no serious work has so far been done on the problem. A typical LEP experiment may need to interface to internal minicomputer buses, plus CAMAC and FASTBUS, some standard microprocessor bus, such as VME, and a local area network. There is a very real problem of how to interconnect everything together in an effective and transportable way, and further work will be necessary soon.

The need for communications systems

Three different types of communication can be identified: (see Fig. 2)

- local communications at an experiment
- CERN wide communications
- wide area communications between the different institutions which make up an experiment team and between these institutes and CERN.

Working group 2 has considered these topics in detail. There have been discussions between LEP groups and WG2, and also cooperation between WG1 and WG2. The comments which follow are therefore somewhat brief relative to the importance of the subject matter, and reference should be made to the WG2 report for more details.

At the level of a single experiment there is a need to interconnect and integrate computing equipment associated with different detectors and sub-systems. This interconnection problem will involve many processors of diverse types, micros, minis and special purpose devices, and a large number of terminals, all of which need to communicate in a coherent way at reasonable cost. Current developments in the field of local area networks (LANs) make this type of technology of considerable relevance and interest to LEP groups. CERN should recommend appropriate solutions which could fit the needs of all experiments. They should be widely available and supported, and follow international standards wherever possible. A basic set of protocols should be available to support task to task communications, remote terminal access and file transfer. However, serious consideration should be given to the standardisation of the higher levels of protocols which are likely to be required. Examples of these types of protocol are:

- control and synchronisation of data acquisition in different processors.
- event sampling.
- remote error reporting.
- remote histrogramming.
- remote control.
- remote CAMAC operations.
- down line loading and control of microprocessors.
- basic house keeping functions for a multiprocessor system.

Existing LANs are in general unsuitable for handling the high rate traffic associated with data acquisition and it is assumed CAMAC and FASTBUS would be used for this purpose. Similarly as LEP event sizes will be large the use of LANs to sample raw data should probably be avoided. When sampling is required a more satisfactory technique would appear to be some form of "spying" directly on the readout buses, as is done now with the present Romulus system. Where many cheap microprocessors need to be interconnected most existing LANs suffer from the disadvantage of a relatively high connection price. The requirement for a cheap medium performance LAN does not appear to have been studied by either the recent microprocessor study groups or WG2, we recommend this problem should be addressed in the very near future. Some practical experience of LAN technology would be very valuable to CERN, as at the moment the organisation is rather short of "hands on" experience in this area.

As in the past at the ISR and SPS communications will be required between the on-line system of a LEP experiment and the machine control computers. The type of interaction required between the two systems may be rather simple; status information is given to an experiment on demand or at regular intervals. Alternatively detailed information may be required from LEP experiments in order to help optimise some aspect of the collider's performance. Precisely what is needed is unclear at the moment and will require further study.

The final requirement for communications at LEP experiments will depend on the detailed disposition of computing equipment and people in the underground areas and in surface buildings.

On-site communications needs are of considerable importance given the enormous geographic extent of LEP. Experimenters are assuming that CERN will provide a laboratory wide communications network to allow access to and from local, experiment based, computers and terminals and the rest of CERN site. Although at the moment remote data recording of large amounts of raw data appears unlikely, it has not yet been totally excluded nor should it be as new developments in technology may make such a strategy viable and attractive in the future. The possible requirement to record data remotely or transfer large event samples between LEP experiments and the main CERN site leads us to recommend that, in addition to medium speed general communications via a general CERN "backbone" network, dedicated high speed links to LEP experiments should be possible (even if not initially implemented).

Several clear facts have emerged about wide area networking. Firstly at least two LEP groups have made very strong statements about the need for this type of facility and all LEP groups recognise its importance. Secondly this need is not just for communications around Europe, but on a truely world-wide basis. The ECFA initiatives in this area seem to have been very fruitful and should be strongly encouraged. The point is not if but how we implement such facilities, and what CERN's role should be.

Effective wide area networking within the high energy physics community would offer many major advantages including

- an active participation of all physicists in LEP experiments.
- the effective use of the computing facilities available to a collaboration.
- efficient data analysis coordination.

We recommend CERN supports and encourages the establishment of this type of service with high priority.

To summarise, communications are one of the most important areas in computing for LEP. They are felt by many to be the most important single topic.

Spares and maintenance for electronic equipment and computers

This is a service provided by CERN which is very much appreciated by outside teams. Our standards must not slip at LEP despite a need to tighten the financial belt. The complexity of modern electronics is such that ever greater demands will be made on our service and maintenance personnel. Where necessary appropriate re-training must be provided. We need to invest still more effort in this area for the future, not less. It will be important, in collaboration with outside laboratories and institutes, to rationalise and publicise which equipment will be supported as a CERN standard and what form this support will take.

Reliability and fault tolerance

This topic has not been studied in any great detail by the working group. However, there are some clear areas where reliability can be and must be improved.

Given the enormous number of detector channels that will be present in future experiments, radically new solutions to the traditional cable and connector problems should be sought. This type of problem as well as that of electronics fabrication and reliability could be eased by the design and development of custom-built integrated circuits for detector electronics.

With electronic systems becoming more and more complex there is a growing need to incorporate features at the design stage which allow self testing and automatic diagnosis and location of faults.

The need for continuing discussions and for collaboration with outside institutes

The working group recognises fully that many of the technical areas discussed in this document will require further study. Definite decisions and commitments by LEP groups may not be possible in the short term. We recommend therefore the dialogue started within WG1 should continue on a long term basis.

The working group, in particular members from outside CERN, concludes that European wide collaboration on a wide variety of topics is necessary if cost effective solutions to problems are to be found. This is particularly true because of the large amount of equipment coming from member countries. At the moment no formal mechanism exists between laboratories to foster such collaboration although we recognise that the role of ECFA in promoting data handling standards has proved very valuable. We therefore strongly recommend that steps be taken by the management to set up a suitable framework for the future.

Trends in technology

This section lists some of the technical areas we feel will be important for the future.

1) The dramatic price drop in the field of customised micro-electronics, VLSI, is of great significance. It will become increasingly cost effective to produce specially designed integrated circuits for such things as the component parts of a detector electronics chain, protocol circuits for interfacing to bus systems and networks, etc.

2) To facilitate the design of VLSI, as well as complex and large printed circuit boards, modern electronics CAD facilities will become ever more essential. CAD should allow a more efficient use of manpower, shorter lead times from design to production and improved documentation. Mechanical CAD may have an important role to play in the mechanical layout of LEP experiments.

3) Personal computers and local area networks will play an ever increasing role in our lives, we however, need to gain practical experience in order to harness them to our needs. Decentralisation of computing risks to bring problems as well as advantages.

4) Wide area networking will more and more allow us to increase the working efficiency of large far flung collaborations of physicists.

5) The number of different buses and networks available will increase and so will the need to couple them together

6) The price of hardware per function will continue to drop, whilst the cost of software, being labour intensive, will increase still further. Standardisation will therefore become ever more important.

The composition of WG1

Chairman R.W. Dobinson

Contacts in LEP groups:

J. Barlow	(DELPHI)
N. Gee	(OPAL)
N. Newman	(L3)
W. von Rueden	(ALEPH)

Other LEP experimenters have also participated.

CERN support Services:

D. Jacobs

P. Scharff-Hansen

H. Verweij

and members of their groups.

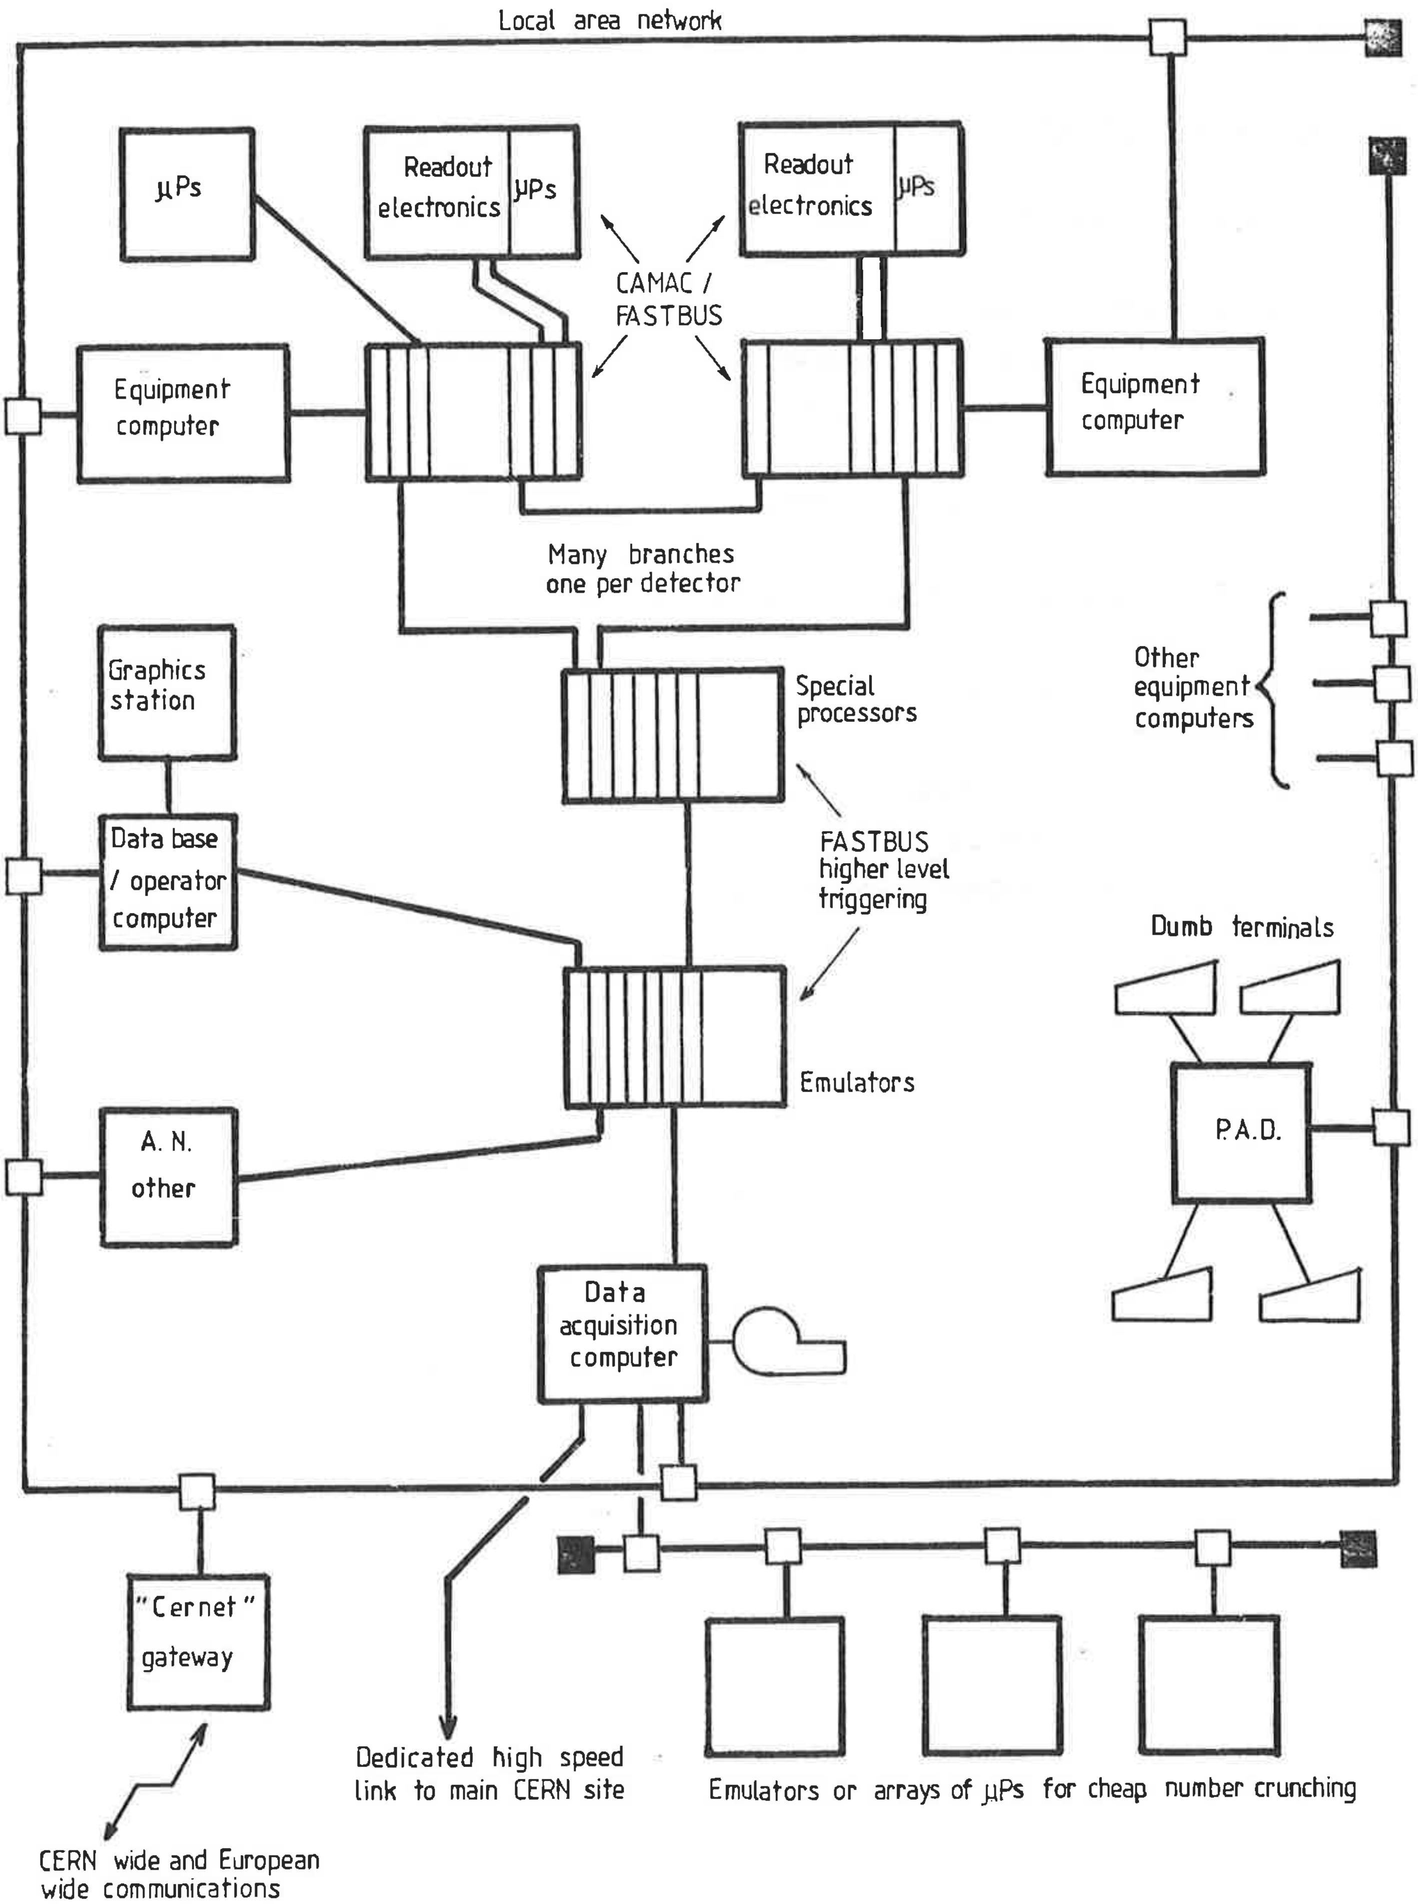

TYPICAL LEP ONLINE SYSTEM

FIG : 1

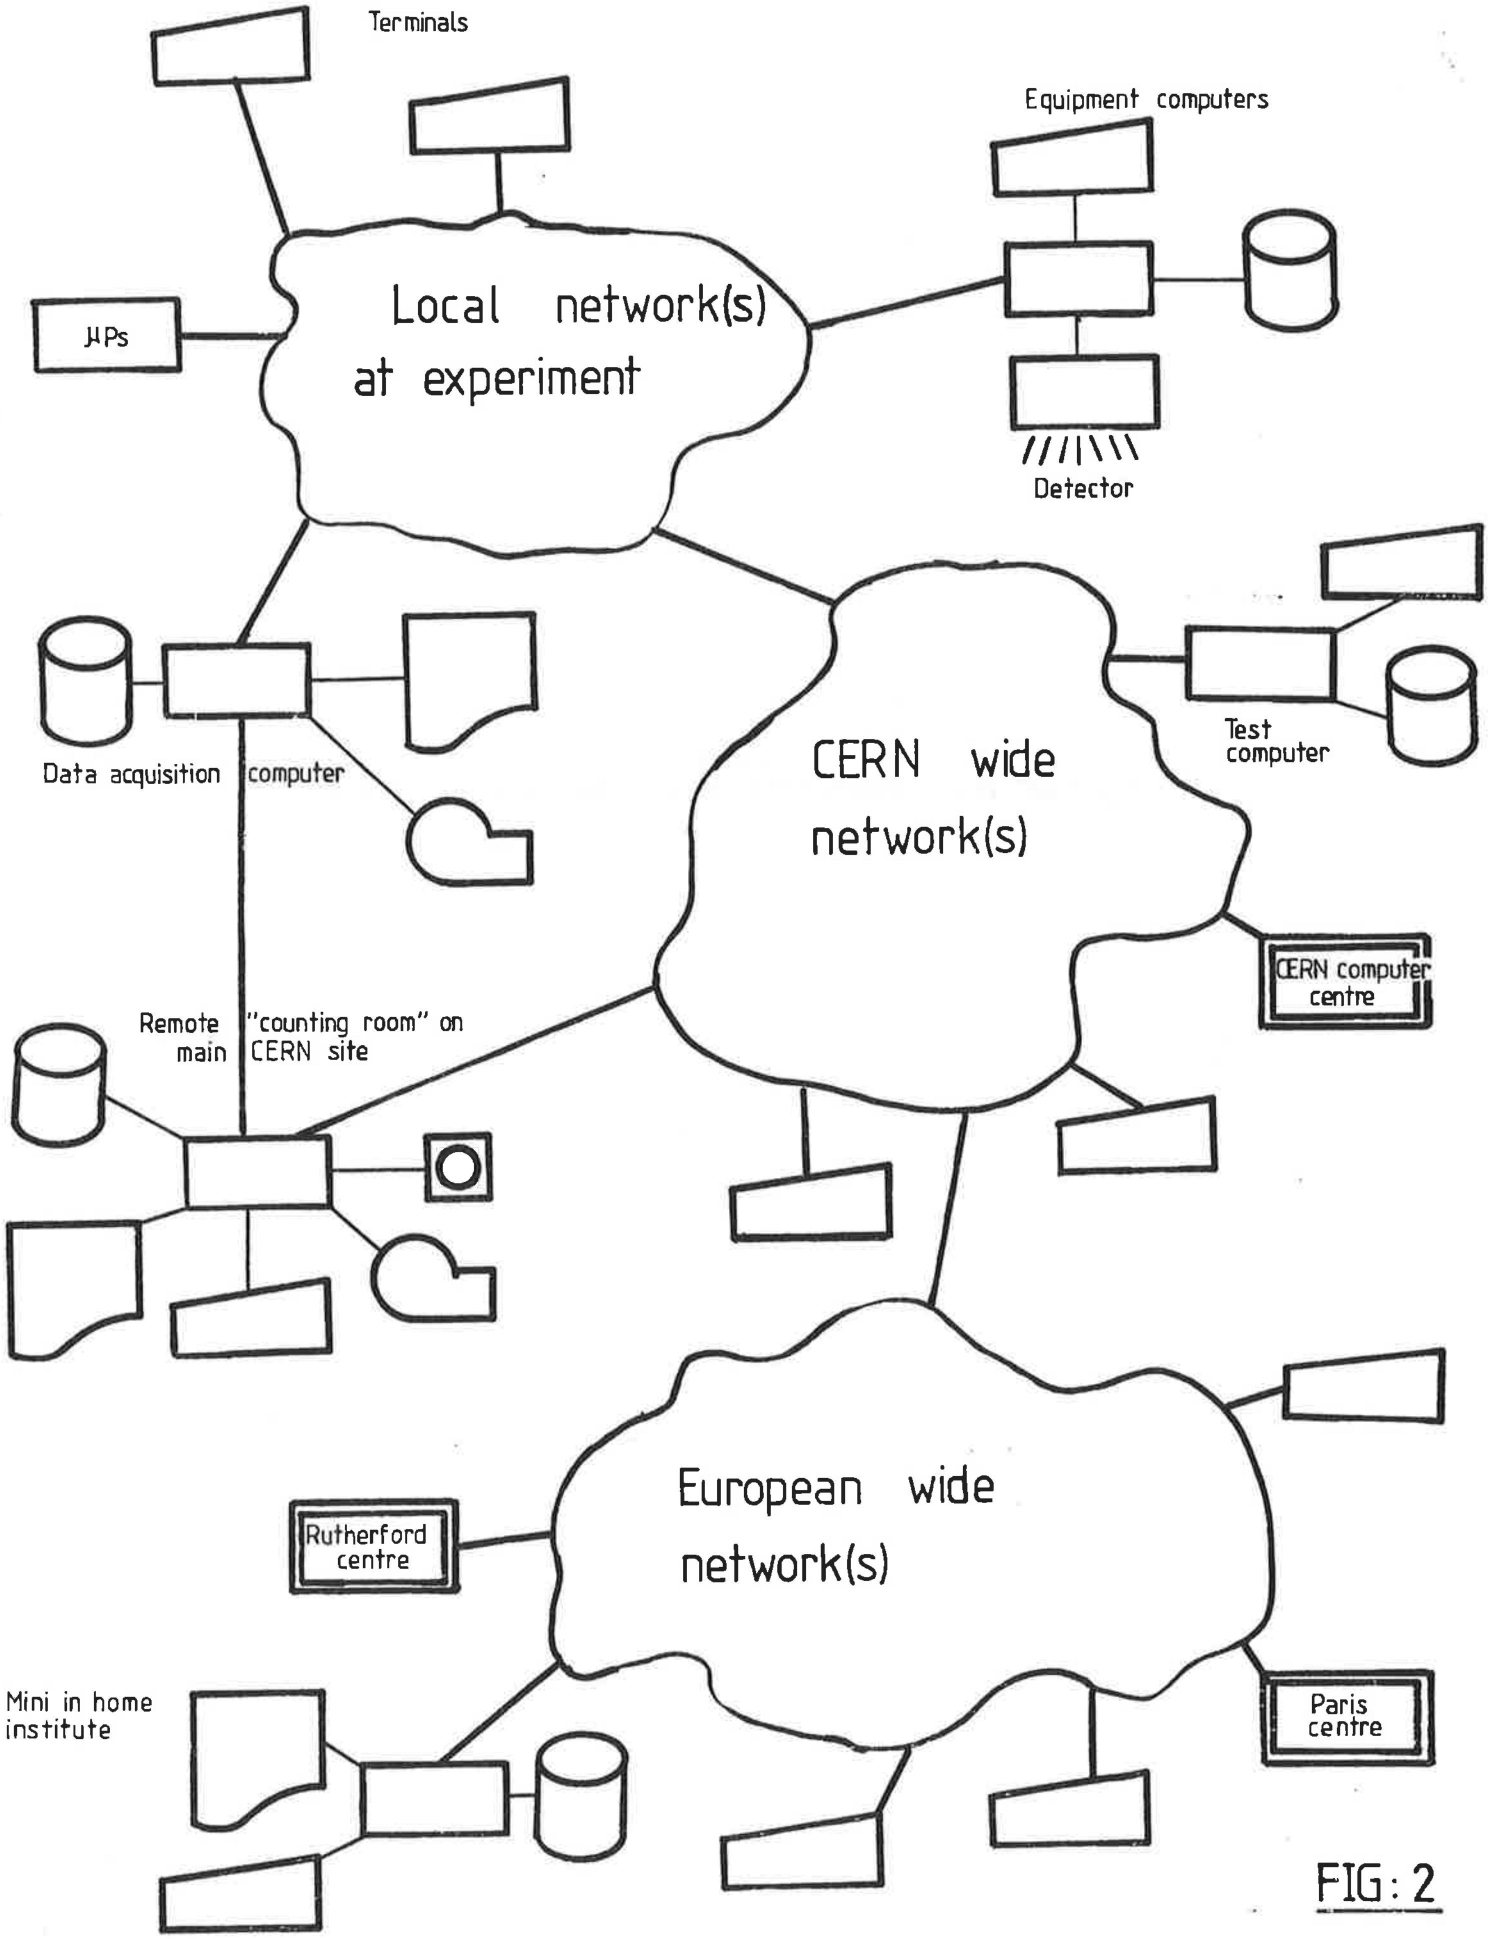

FIG: 2

DIFFERENT LEVELS OF NETWORKING REQUIRED FOR LEP EXPERIMENTS

APPENDIX C
TO THE LEP COMPUTING PLANNING GROUP REPORT

LEP Computing

Planning Group

REPORT

of

Working Group 2

(Networks and Telecommunications)

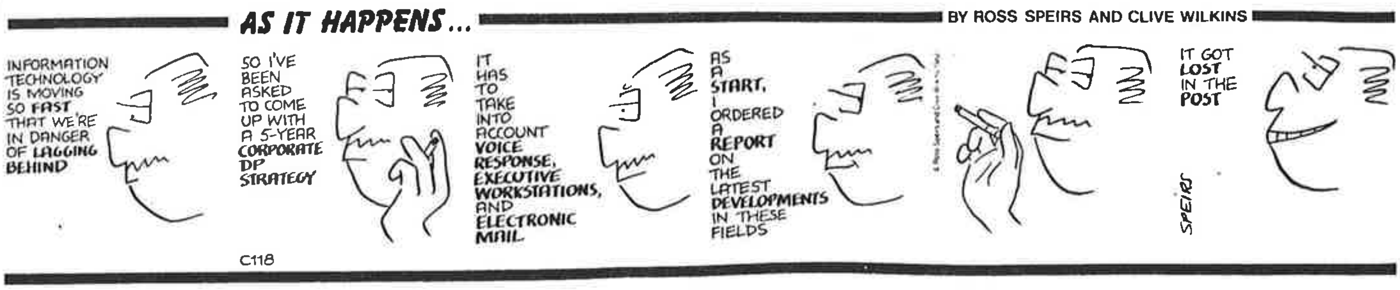

Working Group 2 Final Report

Author: WG2, (ed. J. Gerard)
Date: 22 Dec 82
Version: 13
Status: Final

1. INTRODUCTION

1.1 BACKGROUND

The essential points which have emerged from the working group are:-

1. Data communications is now used by a very large spectrum of CERN staff, spread across all divisions, grades and occupations. The number of users, and types of use, is constantly increasing.

2. The ever-recurring plea from the LEP collaborations is for effective communications facilities, especially including communications to external institutes. Terminal access, electronic mail and reliable file transfer are particularly cited.

3. All installed communications systems eventually recognise the need for connections to other systems, like or unlike. Currently, there is a high degree of interconnection, but on a very ad hoc basis. This has led inevitably to some duplication of effort, and hence a waste of scarce resources.

4. Local Area Networks for inter-computer communication are being actively developed by many manufacturers, and will certainly arrive on the CERN site in the near future. Some of the CERN-developed networks, including the SPS control network, already correspond in function to Local Area Networks.

5. The general purpose network, CERNET, is effectively acting as a necessary backbone for the whole site. However, within two to three years it will be reaching the limits of its capabilities.

6. INDEX, the circuit-switching network for connecting terminals to host computers, will continue to be used to connect simple terminals for some time to come. However, there will be a growing number of requirements for which INDEX alone cannot provide an adequate service.

1.2 PROPOSALS

The working group has, of necessity, considered data communications in its widest sense. It has tried to evaluate present and future requirements, to formulate policy guidelines and to make specific short and long-term proposals. The guidelines and proposals are dealt with in some detail in the next chapter. The following is a brief list of the most important of them:-

1. Future planning for any project involving data communications must be carried out on an inter-divisional basis. This does not prevent the necessary work from being entrusted to a single division. What it does imply, however, is the need for a single, high-level body with overall responsibility for approving plans, and authorising and supervising projects.

2. CERN should play an active part in the specification, adoption and implementation of communications-related standards. This apparently simple requirement may, in fact, involve a considerable amount of effort, but the potential benefits are large. International standards (or likely candidates for future standardisation) should be used wherever possible. The final selection of such protocols for adoption at CERN should be carried out in close collaboration with outside institutes through the ECFA working groups.

3. Continuous monitoring of outside technical developments is essential. Where it is clear that de facto standards are emerging, and that related equipment is likely to arrive and request support, CERN must be ready with a policy defining the conditions under which such support might be given. In cases where such products appear to be of interest to CERN, pilot projects should be set up to install and test limited versions within CERN. The specifications and objectives of these pilot projects should be widely discussed and agreed in advance.

4. Communications facilities between CERN and outside institutes, in Europe or other continents, must be harmonised and improved as fast as possible. Planning should be channelled via ECFA, to which body CERN should guarantee making effort available. The planning must also take into account the effect of permitting access to CERN facilities for these various outside institutes.

5. We recognise that, for some years yet, the majority of simple terminals will continue to use INDEX to access a wide range of computers. On the other hand, we expect that there will be a steady increase in the number of host connections established through packet networks (CERNET, X25 services, etc.). The final goal should be to provide the user with terminal access as part of an integrated communication service, while making the most cost-effective use of the available equipment.

6. Within CERN, the communications architecture should be composed of local area networks, of types approved by CERN, interconnected via a backbone network. Direct connections onto the backbone network should only be permitted in well-defined circumstances.

7. A project should be launched to define and install a high performance backbone network spanning all of CERN, including LEP. The capital cost

is likely to be 2 million Swiss Francs over the next five years, and the effort about 40 man years.

8. It must be guaranteed that the facilities of current systems will continue to be supported and will be smoothly incorporated into future systems.

9. All planned systems must provide generous margins of capacity in order to handle anticipated peak loads easily and must also allow for a continuous increase in the load. A factor of 10 or more over several years is not unreasonable, given that usage of many of CERN's current systems is increasing by factors of 1.5 to 2 each year.

10. In cases where there is a requirement for very high speed, high volume data transfer between two given points, a direct dedicated point-to-point connection should be used. This principle would apply, for example, in the case where the acquisition and the recording of data from an experiment are to be carried out in separate locations.

11. The LEP control network, the future backbone network and any special-purpose high speed links should at least use the same transmission technology, but may be logically separate since they may have very different design goals.

12. CERN should select its standard protocol for each of the required classes of service (file transfer, terminal access, electronic mail, etc.). Such protocols must be chosen during 1983, but some may need to be interim protocols, to be used until international standards are available. CERN should also specify a limited range of other protocols to be supported in each class. CERN should then study the problems involved in writing the necessary gateway software between each standard protocol and the corresponding supported protocols. This work, and the possible ensuing implementations, should be carried out in collaboration with other member-state institutes.

13. The general policy must be to encourage the use of standard hardware and software and to discourage the use of non-standard products. This implies a complete and well-defined policy on support for standard systems, and, for non-standard ones, a validation procedure designed to protect the integrity and efficiency of the overall CERN communication facilities.

14. The users of the CERN communications facilities should be offered standard high-leve service packages, such as a file-transfer utility. These standard packages should exist for a limited number of supported systems, and should meet the basic user needs. Externally-introduced packages may be accepted, subject to verification that they meet CERN standards.

2. GENERAL

The Working Group has continued to operate on two levels. It has created three subgroups; the first concerned with basic transmission hardware, the second with evaluating the set of services required in computer communications networks, and the third with creating implementation models for the mid to late eighties. At the level of the Working Group itself there is liaison with the other two allied Working Groups, with separate Working Groups considering related subjects such as microprocessor support, and with the ECFA subgroup on Links and Networks. This liaison is ensured by an overlap of membership. More details may be found in Appendix 2.

This report tries to extend the period of study into the long term, i.e. the period up until the start-up of LEP. It is already possible to make some proposals over this time scale. However, it must be emphatically stated that the general subject of telecommunications and allied services is vast, complex and on the threshold of revolutionary developments. The Working Group feels that it is only now beginning to understand where many of the real problems lie. It would therefore be unrealistic to expect this document to indicate the precise solutions to all the future problems, or even to recognise what all of the problems will be.

An important fact is that the impetus for identifying requirements and making proposals on how to satisfy them is coming from a wide cross-section of CERN scientific staff. Most of these staff are vitally concerned, whether knowingly or unknowingly, by data communications. LEP accelerator control is clearly dependent upon data communications; what is less obvious but yet very true is that the hundreds of people in the LEP experiment collaborations, inside and outside CERN, will have an absolute need of good communications. Finally, the advent of mini and microcomputers has extended the domain of computers and allied services into all areas of scientific, technical and administrative work. The inevitable conclusion is that a greater proportion of resources must be allotted to the field of data communications.

It is intended that this report should try to set the foundations of a coherent policy for data communications, with the basic assumption that there must be adequate resources made available. In order to build upon those foundations we strongly recommend that there be some kind of continuation of this Working Group, with the task of studying in more detail the problems already recognised and stating in a more precise manner how they should be dealt with.

Many of the aspects of long-term planning are based upon the reports of subgroups 2 (Characterisation of protocol services) and 3 (Future implementation models). Their reports, given as appendices to this report, are based upon a great deal of hard work trying to foresee future needs. Despite all of this hard work there is still a great deal of study yet to be done.

One particular observation of the Working Group and its subgroups is that there are some needs for which an immediate, if temporary, solution is required. In order to try to bring some co-ordination into this short-term planning, Working Group 2 has disbanded its subgroups and is attempting to form Technical Study Groups in the areas where such immediate needs exist. The membership will be at a technical level and will involve end users and potential implementors, not necessarily all from Working Group 2. Overall co-ordination will be done by Working Group 2 until such time as more official management and co-ordination structures can be put into place.

2.1 THE EVOLVING NEED FOR COMPUTER COMMUNICATIONS

Current developments in computing are rapidly changing the access to, and use of computers. The tendency to allow distributed access to powerful number-crunching computers in a computer centre is now being complemented, and in many cases replaced, by a physical distribution of the computing hardware. This will be further accentuated by the increasing use of microprocessors, whether stand-alone or integrated into experimental equipment, intelligent terminals or personal work stations. The need for a new generation of communications equipment is therefore self-evident.

In the context of CERN and its developments over the next few years communications will serve a vital purpose. The areas of importance may be listed as:-

1. The LEP accelerator will cover a much larger geographical area than the SPS accelerator. Control of the accelerator must be very reliable, since it will be a storage ring, and will involve control at a distance of a large amount of equipment and microprocessors. All of these factors increase the need for good and reliable telecommunications across the LEP site.

2. The LEP experiments will be sited far from the two current CERN sites, and will thus need good communications for both data and computer-based message communication. The former will require a high carrying capacity (bandwidth), whilst the latter will call for sophisticated programming effort in order to produce user-friendly systems.

3. The physicists concerned with the LEP experiments will be part of very large collaborations with many home-site bases. Many will be at their home sites for significant fractions of their time. Communications must therefore cover these institutes and allow meaningful participation for these people.

4. The current physicist user community should be able to benefit from the developments done for LEP communications. This will mean extending LEP communications techniques, equipment and services onto the current sites.

5. The CERN community in general is making more use of computers and communication. In this CERN is following a trend already established outside CERN, where computers are being used by all scientific and administrative staff for purposes such as computer-aided design, simulation, management control, office automation and allied applications. The size of CERN means that such applications will of necessity require computer communications.

As a complementary remark, it has been clearly shown that the major computer companies are devoting a very large effort and expense to networking. Digital are basing their developments around Ethernet. CDC have said that their planning for the second half of the decade is based on network architecture, including Ethernet. IBM have already been pushing their Systems Network Architecture software very strongly, and are giving much publicity to their prototype token-passing ring network.

The major problem for an institution like CERN is, however, not only the provision of one (or several) basic communications networks. Rather, it is the effort required to make such networks usable to the customer. In the past this has meant the provision of software packages at various hierarchical levels (process to process communication, access to remote data bases, terminal usage and so on). All of the major software effort will again be required in the future, with extra constraints and requirements.

1. Integration of facilities across the whole of CERN must be considered as necessary from the outset. Recent developments have shown the benefits of interconnecting the SPS network, the PS network, CERNET and OMNET, and of being able to have terminals connect to the IBM either by INDEX or CERNET.

2. The facilities offered by the current system (with the exception of OMNET) will need to be smoothly incorporated into any future systems. In addition there will certainly be some imported network-based facilities to be integrated, such as Ethernets, Cambridge rings or Personal Work Station networks.

3. A point of increasing importance is the communication with external laboratories in Europe, North America and even other continents. The fact that some of these laboratories have already adopted a range of higher-level protocols and applications will inevitably cause extra effort to be necessary.

Our conclusion from all of the above is that the advent of LEP gives CERN an opportunity to create a unified data communications infrastructure and set of services. The effort involved will be considerable in terms of manpower, even if CERN adopts a more rigorous attitude towards standardisation of equipment, software and services than at present. We consider, however, that the effort which would be required if CERN does not try to create such a unification would in practice be even more considerable.

2.2 REVIEW OF PROPOSALS

The proposals of the Working Group vary from specific costed propositions to completely general guidelines. This section deals with the full range of proposals, with the more specific ones being expanded, together with resource requirements, in the next chapter. It is worth noting that the capital expenses involved are relatively less important than for new computer centre equipment, and that some of these expenses may be charged to experimental collaboration budgets. Effort, however, is crucial, both in development, which must be done in a CERN-wide integrated fashion, and maintenance of hardware and software.

2.2.1 Outside Connections

The provision of links to outside institutes is considered of great importance by the LEP collaborations. The collaboration with the ECFA working group is currently quite successful. However, within DD there is more than one group involved. There are also different access methods (IBM or CDC remote job entry stations, incoming or outgoing terminal access via telephone connections, private leased lines, X-25 gateway). The users and potential users of these services would benefit by having a single CERN entity capable of taking an overview of all of the possibilities and needs. It is to be stressed that in the future such users must be treated in a manner comparable with internal CERN users. In particular guidelines on the relative allocation of resources between internal and external users should be formulated, recorded and publicised.

2.2.2 Local Area networks

The field of Local Area Networks is of the utmost importance, because of potential applications in data acquisition computer and work station interconnections. Within this field there is already some effort (on the Cambridge Ring and the ISR-developed UTI-net). However, Ethernet seems certain to be one of the main lines of outside industrial development. CERN should therefore acquire some Ethernet equipment as soon as possible, in order to gain practical experience. An evaluation of the new problems in using LANs and a comparative evaluation of the three LANs, plus ideas on their use in the context of LEP experiment systems, should then be possible. This would seem to fall fairly naturally into the DD Division program, in close collaboration with the ISR personnel.

2.2.3 CERN Backbone Network

The studies on the capacity of CERNET for data communications between computers show that it is likely to be adequate for a few years yet. It is also reasonable to suppose that it can be extended over the LEP site via the Time Division Multiplexing cables to be installed for LEP control. However, the fact that its normal traffic is nearly doubling each year and it is getting large amounts of additional traffic for which it was never intended to cope (168E usage, on-line data recording) implies the necessity for development of a new backbone network for CERN. Some of the specifications depend crucially upon the amount of LEP experiment raw data to be transferred across it. The possibility of a commercially-available system is not to be ruled out, even though it is a specialised high-technology development.

The studies made by subgroup 3, and reported in Appendix 7, suggest that there is as yet no currently available product which can satisfy the requirements. The resultant conclusion is that the current CERNET will need to be developed in order to cope with future demands until such time as a new system can be provided.

2.2.4 Internal CERN links

There should be a minimum number of basic hardware technologies for the communications links across the LEP site. For the immediate future only one such technology is proposed, namely the CCITT G-700 series standards. The necessary tests should be made to ensure that this technology can be used as a carrier for current CERNET and INDEX services. Future usage will include LEP control, a future CERN backbone network, specific dedicated high-speed links and telephone connections. The technical responsibility for installation and maintenance of the necessary hardware should be centred in one group.

More details on this subject may be found in the report of subgroup 1, which appears as Appendix 5 to this report.

2.2.5 Terminal Access

The use of terminals is constantly increasing in CERN, with installation of simple alphanumeric terminals connected onto the INDEX circuit switching system running at around 10 to 15 per month. In most cases, INDEX will provide a connection possibility for these types of terminal for some years yet. However, in the future there is sure to be more diversity of terminals, such as graphics terminals, terminals with incorporated microprocessors, personal work stations and terminals directly connected onto private Local Area Networks. All of these may need to communicate with some of the various host computers around CERN, or even outside. Whilst Gandalf, the makers of INDEX, may be able to meet some of these needs there will still be an increasing requirement for access via CERNET or its successor. The methods of such access (Virtual Terminal Protocols, Triple-X protocols, bridge software in hosts) need a detailed study as part of a general effort on terminal access.

An area of study which may well be of interest for terminal access is the use of the telephone exchange within CERN. Modern developments in private telephone exchanges make it possible for the lines to be used for digital-type traffic. Since all offices are equipped with telephones there is therefore the potential to bring data communications to everyone in CERN at little additional cost, including terminals, direct access to telex, facsimile and even voice mail. By modern standards CERN's telephone system is primitive, and so any study of a replacement is highly important and should involve wide consultations.

2.2.6 Gateways

A major area of effort will be in the provision of gateway services to allow smooth transition between different systems. This might be for conversion from current systems to new ones, between different new ones or for external connections. The computer industry always regards necessity for smooth conversion as a high-priority requirement worthy of considerable effort; CERN must do likewise.

Most of the gateways will be concerned with automatic conversion between the differing, but functionally equivalent, protocols in different systems.

2.2.7 General Organisation

The desire for an integrated approach to computer communications across the CERN site makes it imperative to have some high-level body to survey all computer-based projects liable to require any sort of data communications. Such a body would call upon technical experts, such as this Working Group or a similar entity, and should then be able to make and enforce recommendations.

3. PARTICULAR PROPOSALS : COSTS AND EFFORT

The working group has made a number of particular proposals, which will require resource allocation. The intention in this chapter is to give a better definition of these particular proposals and to quantify the resources needed for them. Note, of course, that the resources necessary to maintain all of CERN's current services are considerable and cannot be significantly reduced until the new services can fully replace the old ones.

3.1 EXTERNAL COMMUNICATIONS

It has been repeatedly stressed in all working groups and in all discussions with physicists that good and usable external communications are absolutely crucial for LEP. The amount of effort to be put into this area should be made commensurate with the increasing importance of such external communications.

The immediate requirements for external communications have been intensively discussed in the ECFA subgroup on links and networks. The conclusion there, which matches the requirements of the physicists, is that there are three obvious and immediate needs.

1. Terminal access, either into or out of CERN, is required both for institutes who are already connected to CERN via leased lines and for institutes having connections to X-25 public networks. The access into CERN should be to virtually any computer connected to an experiment, as well as the computer centre mainframes and special service systems such as data bases and document handling systems.

2. Transfer of source files such as programs, Patchy updates and so on, must be reliable even if not ultra high speed.

3. Mail and documentation systems are crucial to the effective working of LEP collaborations, and must be provided in a generally accessible manner.

> We recommend that there should be some grouping whose members deal with the entire range of connection possibilities between CERN and the external world. In particular, there must be more assigned effort to the X-25 programme of work. A figure of 3 to 4 CERN staff on this latter area is a minimum. This would allow CERN to take an active part in co-ordinating work done at CERN and at the other laboratories.

3.2 LOCAL AREA NETWORK EVALUATION

CERN already has the Cambridge Ring, for which an evaluation effort is in progress. Equally, there is the ISR-developed UTI-Net, a Local Area Network aimed at low cost connections for intelligent devices within a single building and having compatibility with international standards. CERN also may be solicited for a collaboration with the Zurich IBM research Laboratory on their Token Ring. Despite all of this, CERN cannot afford to ignore Ethernet.

> Our specific proposal is to purchase an Ethernet for installation in buildings 513 and 31. A reasonable configuration would be six computer interfaces (2*VAX, 2*PDP/11 and 2*Nord-100), four Personal Work Station interfaces (2*Apollo and 2*Perq), a 4-port terminal interface and eleven transceivers. The cost should then be below 100,000 Swiss Francs. Hardware effort should be not more than one person, whilst software effort would eventually grow to 2 to 3 people, with an active collaboration from DD/OC group. Where possible, software should be obtained from the computer manufacturer. A possible alternative source might be a software house; some software houses are well advanced in the provision of software.

The whole area of Local Area Network development and support needs integrating, so as to avoid duplication and make a meaningful programme of work. The required size for whatever grouping of people deals with LAN development and support, including such things as connection of the Apollo Domain LAN, could very easily grow to 10 or more people (inclusive of hardware and software development and maintenance) as soon as LEP experiments and other users begin practical work with LANs.

3.3 CERN BACKBONE NETWORK

In order to put in perspective the costs and effort which will be required for a backbone network it is worthwhile to consider the equivalent costs for what exists today. Essentially, the current backbone for computer connections consists of a combination of CERNET and OMNET. The elder of these, OMNET, was developed mainly in CERN, and so the direct cost of around 500 000 Swiss francs is certainly exceeded by internal software and hardware development and maintenance costs.

The total capital expense of the CERNET two-phase project from 1975 to 1980 was about 3 million Swiss francs plus 62 man years of effort. The current cost, essentially maintenance and installation of new subscribers, is about 750 000 Swiss francs and 9 man years per year. The portion related to installation of new subscribers, approximately 40% of the yearly capital cost, is now charged directly to the new subscribers.

Subgroup 3 considers that currently available commercial products cannot satisfy the essential requirements for the backbone.

> For this reason, following Subgroup 3, we recommend continued effort to extend the capabilities of CERNET to serve as an initial backbone network. In parallel, starting in 1983, a project should be set up to specify and implement a high performance backbone designed to supersede CERNET from 1986 onwards.

3.4 LEP CONTROL NETWORK

The LEP control network will be used to interconnect the process control multi-microprocessor assemblies and other computers. It will be in the form of a number of interconnected stars. The fundamental property of the network will be the guaranteed delivery of datagrams, together with the ability to assign priority to datagrams.

The highest levels of protocol must follow those of the SPS control network. The eventual objective is an integrated network across the various accelerators.

For the intermediate levels there is the possibility of a choice. Such a choice could facilitate interconnection with the services already available in CERN, as well as with the future data acquisition systems to be constructed by the LEP collaborations. However, the choice must preserve the qualities of service necessary for the applications purely within the scope of LEP control.

The types of interconnection which could be envisaged between the LEP control environment and the rest of CERN could include:-

1. Terminal access from within the LEP control network to the rest of CERN (e.g. Wylbur or Oracle), but not the reverse (i.e. no direct connection of an experimenter's terminal onto the LEP control computers).

2. File transfer to and from other systems, including the experiments. For the particular case of the central computers, file transfer can include job submission and output retrieval. This may need certain restrictions to avoid blocking the control network.

3. Mail and operator messages to and from different sources. This should not cause problems since each normal mail message would typically occupy only one datagram.

> We consider it necessary, in the very near future, to study the question of these intermediate protocols in the context of the proposed LEP network and the future CERN backbone network, in order to ensure a maximum of compatibility.

3.5 A PRACTICAL G-700 SERIES TDM IMPLEMENTATION

The use of G-700 series cables and transmission techniques to carry current CERNET traffic is being tested in a laboratory environment. In order to have a realistic practical environment it is proposed to install and use such a system between building 513 and the North Area.

The arguments for this choice are numerous. It would improve the CERNET communications, which currently connect on cables including repeaters powered independently and susceptible to power problems. It would considerably increase the communications potential between the two sites, which will be very important since both building 513 and the SPS/LEP control room are centres for communication. It would allow tests of different cable types and multiplexing equipment, especially the margins allowable for cable in an accelerator tunnel. In the long term it would be integrated into the complete LEP system.

> There is a technical proposition (document T/W43 of the working group documents) in which the estimates for resources are 80 000 to 100 000 Swiss francs, plus about 1 man year of effort. This is really very small in relation to the long-term benefits involved. We recommend that this limited project be set up as soon as possible, with the collaboration of DD and ISR staff (the latter have already been testing G-700 series equipment for LEP), plus the interest of SPS staff.

3.6 PROTOCOL STANDARDISATION INVOLVEMENT

The process of standardisation of high-level protocols is a very slow process indeed, and one which is always somewhat at odds with the desire of the computer manufacturers. Being centrally sited in Europe, in a logical as well as a physical sense, should allow CERN to exert some influence upon the choice of standards within its field of interest.

The practical aspect of a proposition for more involvement in standards essentially comes back to assignment of effort. It takes time to be involved in the various standardisation efforts in the ISO, the CCITT, ECFA and so on.

However, the potential gain from influencing standards is immense.

> We recommend that for the various standards currently under discussion, namely Transport Service, File Transfer, Remote Job Entry and Virtual Terminal, a person should be assigned to follow and if possible try to influence the standardisation effort. Each of these people must have enough time available plus the travel budget necessary to attend relevant meetings.

For the above-mentioned protocols, plus a Down Line Load protocol, it is the view of subgroup 2 that there should be a defined CERN standard by mid-1983. In addition, there is a need for some mechanism by which to test and validate implementations of this standard, given that some protocol implementations may be imported or privately written. The standard may be only an interim decision made in order to guard against an unco-ordinated proliferation of protocols. When truly international standards become available they should, where possible, be adopted.

It should be noted that, initially at least, these 'standard' protocols will be of use as gateway protocols for existing systems.

3.7 GATEWAY SOFTWARE

In reality this will be the area calling for the most effort of all. Despite the best intentions, CERN will have to live with a variety of products and protocols for any given class of service, and these products will have to be integrated.

The basic premise made here is that for the various classes of service CERN should define a standard. This standard may be replaced at a later date by a 'better' one, i.e. one internationally recognised or technically better. CERN should then implement gateway software for connection between CERN standards and any equivalent but non-standard products which CERN decides to support. By this means, the number of gateways for any given class of service will merely be the number of supported but non-standard products for that service.

Considering the classes of service as the higher levels of the ISO model one can categorise them as Transport Level, File Transfer, Remote Job Entry, Computer Mail and Terminal Access. There is perhaps some overlap between some of these. In each case the number of options is likely to average around four, which leads to a figure of twenty or more gateways. Assuming that the software for any gateway might be written by a suitably qualified person in around 6 months (perhaps optimistic) the effort required will be upwards of 10 man years.

It is to be noted that CERN and the ECFA subgroup have already made a proposal to produce File Transfer gateways between selected protocols. This should considerably increase the understanding of the problems involved in writing such gateways, and it is vital that CERN take an active part in the proposal.

> We recommend the above policy for the provision of gateways. We wish to see adequate effort assigned and co-ordinated in an effective fashion. The co-ordination should include a close involvement with ECFA.

4. FUTURE SCENARIO AND PREDICTIONS

This chapter is a review of some of the history of data communications inside CERN, together with the predicted needs of the various user communities, so as to see how to move smoothly into the future.

For the future predictions, what is attempted is to look at the short, medium and long term, corresponding to 1983, 1985 and 1988 respectively. In essence the needs of 1983 must be satisfied by what is available now. However, the facilities necessary for the second half of this decade must begin to be developed next year, in order that they are tested, installable and maintainable in a service environment from 1985 onwards.

In practice, we expect that the specifications for what future facilities will be necessary may change with time and experience. Thus, it is vital to design any system as open-ended as possible. What is considered as a luxury today may become a practical necessity some time in the future.

4.1 THE LESSONS OF HISTORY

A look at past history can very often give a good indication of future trends and requirements.

4.1.1 Continuity

It is clear that CERN is now too large to be able to make a clean jump into new techniques in computers and communications. There is such a large amount of vested effort in the current systems that a smooth continuation is necessary. DD division did actually make a clean termination of the FOCUS system for file editing, data-sample collection and off-line job submission some years ago. However, the number of physicists and experiments was so small that the effort could be made available to convert them to the corresponding CERNET facilities. Any new facilities, such as LEP control system, new general network or terminal access mechanism cannot simply replace the current facilities from one day to the next: they must be smoothly integrated.

4.1.2 Expandability

The growth in the use of computers and communication facilities is equally something which will carry on in the future. Experiments are becoming more complex and requiring more sophisticated facilities and more processing power. Computing is becoming more of a distributed nature, involving more communications across wider areas, and it is reaching into all levels of human endeavour. One therefore cannot artificially limit this growth without seriously inhibiting the overall efficiency of CERN.

Growth in the facilities offered by a computer centre, logically being controlled centrally even if physically distributed, is properly the concern of Working Group 3. There is an obvious cost factor in that the sum total of expenditure is relatively large. The influence of such arbitrary decisions as the need to do two thirds of the computing outside CERN, and the exact interpretation of such decisions, is considerable. History shows that up to the level of final bulk data processing the experimenters prefer to work with CERN facilities. If they cannot, then there is a corresponding extra demand for communications facilities.

4.1.3 Integration of Facilities

Growth in communications facilities is highly dependent upon human resources. It is particularly important here to look at the history of what these resources have done. The following is a brief, incomplete list.

1. At the very start of the operation of CERNET, the development of gateway software and hardware allowed interconnection of CERNET and the existing OMNET network of PDP/11 computers. This gave all of the OMNET users access to CERNET, and hence the main computers.

2. The leased lines connecting CERN to the other European high energy physics laboratories have been adapted to a multitude of purposes, often executing in parallel due to the use of multiplexing equipment. These purposes include the attachment of IBM and CDC remote stations and the connection, via INDEX, to external INDEX equipment. Note that in each case a considerable effort is required to ensure smooth integration and continuation of these services without perturbing internal CERN services.

3. The provision of suitable software has meant that the CDC and IBM mainframes in the computer centre can communicate with each other for file transfer, job submission and output retrieval. Again, the effort involved has been considerable but the results are very much appreciated by the users.

4. The SPS control network has been taken as a basis for the PS control network, such that the two networks can communicate with each other.

5. Both the SPS and the PS network have been connected to CERNET. In addition to the general use of these connections to allow dedicated PS or SPS terminals to access the IBM Wylbur service an SPS computer regularly sends to the IBM information files concerning beam parameters. These are now required for processing of UA1 data.

6. The CERNET X-25 gateway has a connection to INDEX for the purposes of terminal access to and from external institutes.

All of the above shows that inter-system communications have always proved to be necessary in the end. The stage has come where almost any system can somehow access any other one in CERN. However, this has been done in a haphazard manner and resulted in a very inhomogeneous entity. The lesson of history is that CERN must make the effort to integrate communications on a CERN-wide scale. The effort will be considerable, but much less that that which would inevitably be spent if CERN does not adopt such an approach from the beginning.

4.1.4 Timescales

Another lesson of history is concerned with the elapsed time between development of facilities and their acceptance. The experience of CERNET, begun as a project in 1975, is that only by about 1980 did it begin to be seen as an accepted part of a data acquisition system. Other systems, such as INDEX, not reliant upon software effort, have taken much less time to become accepted. Accelerator control systems are accepted immediately because the accelerator and the control system are developed together, but they are worked upon well in advance of accelerator availability. The lesson is thus that for a time scale going up to the late eighties work must begin now.

4.2 USER NEEDS

It is of clear importance to try to foresee what will be the needs of the users of data communications. This subject has been treated by subgroup 2, whose report is included as appendix 6 to this report. Note, of course, that there are qualitative (i.e. which services) and quantitative (how much) needs. Subgroup 2 has attempted to identify the services needed (at the level of network protocols) and to characterise them quantitatively.

4.2.1 Controls

The accelerator control is typical of what one can call control applications in general. LEP experiments equally may have control requirements, as may SB services. The typical requirement is for a variety of protocol services with fast response to short commands, plus a very high network reliability. A proper design aim is compatibility, such that there can be two-way interchange of information between LEP control and LEP experiments.

4.2.2 Data Acquisition

Experiments have a large variety of communications requirements. Many of these arise from the fact that their data acquisition systems are now very complex, with many heterogeneous mini and micro computers. The general principle will continue to be that the treatment, sharing and possibly recording of raw data is internal to the local computers using dedicated high-speed data paths (e.g. FASTBUS), but that these local computers require generalised support from elsewhere. Compared to the present situation there will be an increased emphasis on a fully integrated, reliable and supported communications system with task-to-task communication, terminal access, file transfer, job control and microprocessor support. In particular, the physicists must be able to monitor the data acquisition system without being in the counting room or even at CERN.

The proposed solution to the experiments' general communications needs is a Local Area Network (LAN) interconnecting the various minis of each experiment. Ideally, such a LAN should also allow connection of numerous micros. However, the current high connection costs tend rather towards connecting such micros onto a cheap bus, which in turn has a connection onto the LAN.

The exact role of such a network will evolve from experience and testing. The current range of LAN speeds do not allow their use for the transfer of raw data, whilst their limited geographical spread poses a problem for access from most areas of CERN. Indeed, if remote data recording is ever adopted, with a substantial part of the experiments' dedicated computing power many kilometres from the detectors, some form of split LAN with a transparent connection between its two parts would be needed.

4.2.3 Connections beyond the Experiment

Connections between the experiment's LAN and the rest of CERN (terminals and computers), and transparent connections from the LAN and CERN to the home institutes (Europe, N. America, USSR, Asia) are essential, with a special emphasis on an integrated electronic mail facility.

On the particular point of the manipulation of LEP experiments' pre-processed data, the current estimates are of data rates of 2 megabits per second. It is unlikely that any general purpose network can regularly handle such rates. If it is felt necessary, even for short periods, to sustain such rates over CERN public facilities then the connection should be direct from the source to the required destination. In other words, there should be a dedicated fibre or cable of high capacity, say 34 megabit per second to allow a suitable margin. The LEP collaborations at present envisage local (i.e. under the Jura) data recording and event processing, with remote interactive access. However, practical reasons, and technical evolution, may change this view in due course. In this case, the problem of a geographically split LAN, as mentioned above, would arise.

For the communications external to CERN, the user requirements do not involve very high speed transfer of raw data. It is felt to be more important to have guaranteed reliability. This is particularly obvious if program files are being transferred. The occasional transfer of a large amount of data on an overnight basis should be made possible.

4.2.4 Support Services

A very large community of system (hardware or software) developers and maintainers needs access to the facilities provided by the central software support teams. Cross-support (compilers for high level languages, assemblers, link loaders) are an obvious need, as are access to mass store, fancy output devices (plotters, laser printers, phototypesetting devices) and brute number-crunching capacity.

4.2.5 Personal Support Services

The non-scientific community of CERN also has requirements. The word processors at present in use in various divisions require some means of intercommunication, possibly using the IBM file store as an intermediate medium.

4.2.6 Overall Requirements

It is implicit in all the above that equipment from a wide variety of computer manufacturers must be interconnected flexibly and transparently, i.e. a heterogeneous or 'open' communications system is needed, in general accordance with the ISO 'Open Systems Interconnection Reference Model'.

The above is a very brief list of user needs. The report of Subgroup 2 gives more detail on user needs, domains of application, services and characteristics of services in each domain. It is worth noting that the ranges of protocol services needed for the various domains are not particularly different. Unfortunately, some domains (e.g. LEP control) have urgent decision deadlines and so can probably not wait for a CERN or industry LAN standard to emerge. Nevertheless, Subgroup 2 has recommended a four-point programme for adoption of standard protocols to cover almost the entire range of services needed on the CERN sites and connections to outside. Such protocols would initially be used as gateway standards, and progressively introduced for internal use in new LAN implementations.

4.3 THE SHORT-TERM SCENARIO (1983)

4.3.1 Start of LEP

1983 will clearly be a year of preparation. The services offered will not be very different from those of 1982. Nonetheless it will be important in a number of ways.

The construction work for the LEP accelerator will be starting. This will require computer communications from the outset. The very first requirement will be for telephone communications. Immediately after that will come the need for simple terminals at the excavation areas accessing a reporting and message filing computer service. This may need modems or acoustic couplers until proper lines are available.

4.3.2 Early Support for Collaborations

The experimental collaborations will still be at a very early stage of thinking. They will certainly be aware of the developments in computing and communications, and should welcome informative discussion on the application of these developments. In particular they will have heard much about Personal Work Stations and Local Area Networks. CERN must have in-house experience of these in order to give advice which is correct in the CERN context.

The members of the collaborations will be geographically separated by large distances for much of the time. The possibility of wide area communications across Europe and North America will exist with the interconnection of the PTT networks. There will remain a large amount of work to be done in order to offer practical services, in particular terminal access, computer mail and file transfer.

4.3.3 Extension of System Support

Internal to CERN there will be a continuing growth in communications requirements. Different computer mail systems will generate requests for their interconnection. The growth in the number of microprocessors will call for an efficient support system including cross-compilers, cross-linking, down-line loading, on-line debugging and up-line loading.

The number of terminals will continue to grow at the current rate (around 10 to 15 installations per month), with a corresponding expansion of INDEX port and terminal interfaces. CERN should start to see the more sophisticated terminals incorporating bit-map displays and microprocessors. The demand for fully interactive systems may well rise in priority. CERN should therefore be reviewing its future policy for terminals and generalised terminal access to multiple host systems, to see where INDEX-style circuit-switching is better or worse than network-based virtual circuits.

In general it can be said that 1983 needs to be a year for gaining experience with new equipment and techniques, and evaluating their impact upon CERN's future.

4.4 THE MEDIUM-TERM SCENARIO (1985)

It should be possible to foresee with reasonable accuracy where CERN would like to be in 1985. This choice of date is such that CERN can make decisions now, or early in 1983, and begin to see the results of the decisions then.

4.4.1 Start of LEP Network

For at least part of the LEP site there should be a cable network. This should include in its functionality the LEP accelerator control needs, the carrying of CERNET and INDEX, and any other needs such as telephones. Capacity should also foresee physics raw data, and should thus be potentially high.

4.4.2 Telephone System studies

For the telephone system it should be known if, and when, CERN might install a modern digital exchange, and also the extent to which such an exchange might be used for other services such as terminals, facsimile, remote slow speed printing.

4.4.3 Extension of Terminals

The number of simple terminals installed in CERN should be approaching 1500. The majority will still be relatively simple alphanumeric display terminals. However, there will be a significant number of such simple terminals physically attached to particular computers or local area networks, but wishing on occasions to access general services. In addition there will be a number of microprocessor-incorporated terminals and terminals incorporating point-addressable bit-map displays.

4.4.4 Reassessment of CERNET and INDEX

CERNET will be heavily loaded, and CERN should be in the test phase for its successor. INDEX will still be used for a very large number of terminals in CERN. There will be an increased demand for terminal connections for which INDEX is not immediately suitable, and alternative systems for terminal support will be increasingly used.

4.4.5 Reassessment of the Computer Centre

The evolution of the computer centre could have a considerable bearing on the needs for communications. If the CDC 7600 is replaced by IBM-compatible equipment then the inter-mainframe communications may be supplied by the manufacturer. However, if it is replaced by another CDC (with a different operating system) or another manufacturer then the inter-mainframe communications will pose an additional conversion problem. The choice may be between current CERNET (a stopgap), CERNET successor or a dedicated product such as Hyperchannel.

Another factor of great importance for communications is the role of the computer centre regarding file storage in large quantities. An increase in the size of the IBM Mass Store seems difficult to avoid. What is more questionable is the effect of possibly accessing this Mass Store, or any other medium (video disks?) via a front-end File Server machine. CERN should have made its plans in this area, and there should be a provisional decision about usage by LEP experiments for storage of raw data.

4.4.6 Off-line Emulators

The future use of off-line emulators is unclear. By 1985 the next generation of emulators (the 3081E) should be available. CERN should have decided if they are to be used in an off-line role. If so then they should be supported computer centre services, with an adequate high-speed connection to the computer centre file system, if possible other than CERNET.

4.4.7 Initial Inter-LAN Connections

Internal to CERN there will certainly be a number of Local Area Networks. There will be options here (Ethernet, Cambridge Ring, IBM Token Ring?, others?), and CERN will have a support policy for them. Their interconnection and the services provided and required are likely to have been understood, even if the ultimate goal of completely transparent interconnection is not achieved.

4.4.8 Expansion of External Connections

External to CERN it will be possible to access any associated laboratory, institute or university in Europe or elsewhere. Terminal access should be no problem, whether into or out of CERN. Inter-system mail equally should be provided. File transfer will be possible, although opinions differ as to whether there will be a generally accepted file transfer protocol. Even if there is, CERN may not have completely converted to it.

4.5 THE LONG-TERM SCENARIO(1988)

The great unknown. The complete accelerator control should be set up. The LEP experiments also. There should be lots of centrally-controlled computer services (and CPU power) and also distributed computer power not necessarily centrally controlled; how to use and access it. The CERNET successor should be working and integrated. Interconnection of anything to anything should be possible at little effort (provided CERN has standardised on equipment, protocols, software and provided all necessary gateways).

Although some facts about 1988 are fairly certain there must be some doubts about the ability to crystal-gaze so far in advance. Any planning must therefore be as flexible as possible.

4.5.1 LEP Operation

The LEP accelerator is due to have been constructed and in preliminary operation. The control network will be operational. There should be a surface network across much of the LEP site, plus a coaxial cable around the tunnel. The LEP control network will use both types of cable, and should have at least a certain amount of redundancy and backup possibilities. There may also be a coaxial cable circuit around the LEP ring, possibly used for timing signals or CATV-type technology (e.g. TV transmission).

4.5.2 Computing Resources

The physical distribution of computing resources will have continued, leading to a situation in which computing capacity is not necessarily confined to the computer centre, even if expensive particular peripherals are so confined. The physicist should have access to this capacity in an integrated fashion regardless of where he happens to be.

4.5.3 Computer Communications

The successor to CERNET should be standard and available all over CERN. It should provide for virtually all computer to computer communication on and off site. This will include uses in administration, finance, budgeting and other areas not directly relevant to physics. It should also allow for terminal traffic, although it is probable that INDEX will still be in widespread use.

4.5.4 Protocol Standardisation

Standard protocols should finally have been defined for such mundane applications as inter-computer file transfer, job submission and output retrieval, and there might even be available from the computer manufacturers software allowing dialogue with computers of different manufacturers.

4.5.5 Overall Objectives

The objective for the 1988 time frame should be to have integrated communications. This does not necessarily imply that everything connected with computing will be simpler than now, but simply that CERN will have contained the growth in complexity.

CONTENTS

Terms of Reference
Working Group II: Data Communications and Networks

Author: J. Gerard
Date: 11 Jun 82
Version: 1
Status: Final

1. To assess and attempt to quantify the possible requirements for telecommunications at CERN for the next decade, with particular attention to the requirements for LEP.

2. To match the requirements to the various technologies available now or in the reasonably near future, in order to decide on how to provide the infrastructure for CERN-wide multipurpose Communications.

3. To define the methods and interfaces which may be suitable for attachment of a wide variety of user equipment.

4. To give careful consideration to communications with outside Institutes (High-energy physics laboratories, Universities etc.) in Europe and elsewhere. In close collaboration with ECFA, to classify and encourage the adoption of international standards.

5. In general, to recommend a consistent CERN strategy for telecommunications. Such a strategy should take into account the necessity for general guidelines, standards and clearly-defined projects.

6. To make recommendations as to how the various requirements could be satisfied, taking into account the desire for economic solutions, the uncertainty factor of the requirements and the need to allow for expansion possibilities. These recommendations may take the form of alternative options.

7. To specify the resources required, in terms of capital cost, operational expense and manpower, to develop, install and maintain any proposed systems, together with the timetable for these stages. Smooth transition from the current situation is an important factor.

Notes concerning the terms of reference

a) The number of services likely to require some form of digital data communications is large, and the types of service diverse. Amongst the more obvious are accelerator control, experimental support, microprocessor support, office automation, personal work stations and terminals. In addition there is the area of voice and image transmission (e.g. teleconferencing) which might or might not be digital. The term 'Telecommunications' is thus used as a generic term to cover this whole area.

b) Given the wide variety of potential services and the associated requirements for quality of service, it is unclear whether one particular product can provide all necessary services. However, the goal should be to minimise the number of different technologies as far as is consistent with a functionally adequate product. In the particular field of Local Area Networks (LANs) there are currently several different competing technologies, and it is unlikely that de facto standards will emerge for quite some time yet. It is only realistic, therefore, to assume that a wide variety of such LANs will be brought onto the CERN site, and that requests will then be made for connections to other LANs, either similar or different, inside or outside CERN. Such potential anarchy can only be controlled by a strict definition of acceptable hardware and software interfaces for different levels of interconnection.

c) Collaboration with the other working groups and with outside groups, in particular ECFA, is absolutely essential. Within this working group an effective inter-divisional collaboration is very necessary.

d) The working group should use the ISO Open Systems Interconnection model as a basis for the definition of the various layers of services to be offered. These will range from basic hardware to end-user software products. It would seem appropriate that the working group form small working sub-groups corresponding to these layers, where each sub-group is run by a member of the working group but may include specialists not necessarily in the working group.

e) Initial recommendations must be available for Autumn 1982, which may limit the extent to which the working group can propose specific solutions. For the longer term, there should be some consensus as to the way in which further recommendations are to be made and reviewed in the light of experience.

Working Group membership.

All of the Planning Group are automatically accepted for this Working Group. This is essential for maintaining good communication across the groups. The extent to which any particular person of the planning Group may have both the time and the desire to participate actively has to be clarified.

The active part of the Working Group consists of persons who, in general, are not overly burdened with administrative affairs. It is my basic assumption that these persons will have the time and desire to make active contributions. The list is:-

J.M.Gerard (chairman)
C.J.Piney (secretary)
J.Altaber
B.Carpenter
F.Fluckiger
C.R.Parker
D.Samyn
P.Wolstenholme

Finally, in view of the importance of good interworking with ECFA, I have invited D.M.Sendall to join us.

APPENDIX 3

List of WG2 Documentation

Author: J. M. Gerard
Date: 27 Jan 83
Version: 2
Status: Current

List of T/W of DOCP documents on 27 Jan 83

Proposed Program of Work for CERN Pilot Cambridge Ring
V1 : M. Hine,B. Segal,D. Wiegandt,G. Lee,I. Willers (T/W28)
CERNET Node Efficiency Measurements
V1 : J.M.Gerard (T/W29)
Draft note on Short-term applications of the Cernet 68K module for OC and SW
V1 : J.Anthonioz-Blanc,S.Brobecker,J.Joosten,D.Samyn (T/W30)
Invitation to WG2SG3 meeting of 14 Sept
V1 : C. Piney (T/W31)
Notes on a possible Ethernet project definition
V1 : D. Sendall (T/W32)
Summary of WG2 Subgroup 2 Results (Characterisation of protocol services)
V4 : B.Carpenter,D.Samyn,for SG2 (T/W33)
Essential Requirements for all Implementation Models
V4 : R. Parker,C. Piney,M. Sendall,I. Willers (T/W34)
Criteria on which to Evaluate Models of Communication at Cern
V1 : WG2SG3,(ed. C. Piney) (T/W35)
W.G.2 Study Group 1 Interim Report Suitability of CCITT G-700 TDM Systems for CE
V8 : P. Wolstenholme (T/W37)
Rough cost estimates for Omnet.
V1 : D.M.Sendall (T/W38)
Survey of Data Communication facilities at Cern
V1 : (ed. C. Piney) (T/W39)
Implementation Models for a Future Communications Programme at Cern
V11: WG2 SG3,C. Piney (T/W40)
Aims of standardisation in local communication
V1 : D. Sendall (T/W41)
How to enter WG2 minutes
V1 : C. Piney (T/W42)
Possibilities for Installation of TDM Links to North Area
V1 : P. Wolstenholme (T/W43)
Brief note on DEC communication products
V1 : D.M.Sendall (T/W44)
Short note on visit by C. C. Maples (LBL)
V2 : C. Piney (T/W45)
Notes on possible use of DECNET at CERN.
V1 : D.M.Sendall (T/W46)
Preliminary Proposal for the Data Link Unit for LEP
: J. Altaber,M. C. Crowley-Milling (T/W47)
Tabular results of the SG2 survey
V1 : D. Samyn (T/W48)
Aims of the UTI-NET Project
V2 : R. Parker (T/W49)
Comments on RAL suggestion for X25 connection to some VAXs
V1 : D. Sendall (T/W50)
UTI-NET: Status and Plans
V2 : R. Parker (T/W51)
MIL 1553-B
V1 : J.Altaber SPS (T/W52)
Communications in the experiment - some issues
V1 : D.M.Sendall (T/W53)

List of T/G of DOCP documents on 27 Jan 83

Memorandum on Data Networks
: M.C. Crowley-Milling,P.G. Innocenti (T/G1)
Timescale for LEP data networks
: M.C. Crowley-Milling (T/G3)
Chart of LEP Timescales
: (T/G4)
ESONE and Local Area Networks
: A.C. Peatfield (T/G5)
Terms of Reference Working Group II: Data Communications and Networks
V1 : J. Gerard (T/G6)
Measurements on the Control Network of the SPS
: J. Altaber,J.-P. Jeanneret (T/G9)
Trip Report USA 1981
V1 : J.M. Gerard,C.J. Piney (T/G10)
Telecommunications Agreements with Local PTT Authorities
: P. Wolstenholme (T/G11)
Memo on Relations with the Swiss and French PTTs
: M. Hine (T/G12)
Price List for Cambridge Ring Products
: Polynet/Logica (T/G13)
Working Group 2 Progress Report
V2 : WG2,(ed. J. Gerard, C. Piney) (T/G14)
OC meeting on ETHERNET pilot project
V1 : D. Samyn (T/G15)
UTI-NET: An experimental, low-cost local area data network: A Preliminary Report
: C.R.C.B. Parker (T/G17)
Networks for High-Energy Physics An Interim Report on Wide-Area Communications
V1 : ECFA WG on Data Processing Standards SG 5 (T/G18)
List of Members of ECFA WG on Data-processing Standards, SG5
V1 : ECFA (T/G19)
Working Group 2 Final Report
V13: WG2,(ed. J. Gerard) (T/G20)
Minutes of Network Communication Meeting Report on Cernet 68K module and Intermediate speed communications
: G. Macleod (T/G21)
Memo on Future PTT Telecommunications
V1 : H. Hine (T/G22)
Data Networks and the LEP Project
: M.C. Crowley-Milling (T/G23)
CERNET ACC Status, Configurations and further development
: F. Fabiani (T/G25)
Comments on the Report of the Office Automation Panel
V1 : J. M. Gerard (T/G26)
Propositions for a DD Ethernet
V1 : J. M. Gerard (T/G27)
Strategies for File Transfer between Personal Work Stations and Host Computers
: B. M. Segal (T/G28)
Report in a visit to the Direction Operationnelle de Telecommunications d'Annecy

List of T/M of DOCP documents on 27 Jan 83

Minutes of WGII Meeting of 25 May 82 V2 : C. Piney	(T/M1)
Minutes of WGII meeting of 1st June 82 V2 : C. Piney	(T/M2)
Minutes of WGII meeting of 8 June 82 V2 : C. Piney	(T/M3)
Minutes of WGII meeting of 16 June 82 V2 : C. Piney	(T/M4)
Minutes of WGII meeting of 22 June 82 V3 : C. Piney	(T/M5)
Minutes of WGII meeting of 29 June 82 V2 : C. Piney	(T/M6)
Minutes of WGII meeting of 6 July 82 V2 : C. Piney	(T/M7)
Minutes of WGII meeting of 20 July 82 V1 : C. Piney	(T/M8)
Minutes of WGII meeting of 3 August 82 V2 : C. Piney	(T/M9)
Minutes of WGII meeting of 17 August 82 V1 : C. Piney	(T/M10)
Minutes of WGII meeting of 31 August 82 V2 : C. Piney	(T/M11)
Minutes of WGII meeting of 7 September 82 V1 : C. Piney	(T/M12)
Minutes of WGII meeting of 14 September 82 V1 : C. Piney	(T/M13)
Minutes of WGII meeting of 21 September 82 V1 : C. Piney	(T/M14)
Minutes of WGII meeting of 28 September 82 V1 : C. Piney	(T/M15)
Minutes of WGII meeting of 5 October 82 V1 : C. Piney	(T/M16)
Minutes of WGII meeting of 12 October 82 V1 : J.M.Gerard	(T/M17)
Minutes of WGII meeting of 20 October 82 V1 : C. Piney	(T/M18)
Minutes of WGII meeting of 26 October 1982 V1 : J. M. Gerard	(T/M19)
Minutes of WGII meeting of 2 November 82 V1 : C. Piney	(T/M20)
Minutes of WGII meeting of 9 November 82 V1 : C. Piney	(T/M21)
Minutes of WGII meeting of 16 November 82 V1 : C. Piney	(T/M22)
Minutes of WGII meeting of 30 November 82 V1 : C. Piney	(T/M23)
Minutes of WGII meeting of 7 December 82 V1 : C. Piney	(T/M24)
Minutes of WGII meeting of 13 December 82 V1 : C. Piney	(T/M25)

APPENDIX 4

Technical details on transmission products

Author: J. M. Gerard
Date: 8 Dec 82
Version: 1
Status: Current

This document is essentially technical in content. As such, it should provide a background for the reader having a particularly detailed interest in the WG2 final report.

The first section is intended to give an overview of the basic techniques for data transmission. The CCITT G-700 series is of particular interest, but other techniques could also be of interest in the long term.

The plethora of products appearing on the market makes it difficult to review them all. In addition, one must distinguish what are essentially tools from what are complete user-oriented services. The former are described in the second section; the latter in the third.

1. BASIC TRANSMISSION TECHNIQUES

Subgroup 1 of Working Group 2 was formed to look into the suitability of using CCITT G-700 series Time Division Multiplexing systems for the cabling of the LEP site, with the particular requirement of assessing their suitability for extending the current CERNET and INDEX services across LEP. Their Interim Report is included as appendix 5 to this report.

The transmission systems concerned inherently provide point to point links with speed options 2, 8.5, 34 or 140 megabits/second. Multiplexing equipment is available to split the 2 megabit links into units of 64 kilobits, the standard transmission speed for digital telephone circuits, and to split the higher speeds (8.5, 34 or 140 megabits) into four of the corresponding next lower speed.

The physical medium can be twisted pairs, coaxial cable or optic fibre. It is not yet clear under what circumstances, and over what distances, the twisted pairs currently in widespread use on the CERN site may be used. In general one should use coaxial cable at 8.5 megabits or higher, with fibre optics an option at 34 megabits or higher except in the LEP tunnel. The cost of the cable can vary considerably with the quality and characteristics demanded. Of course, the ultimate cost of installing cable of insufficient quality could be considerable: CERN really needs to gain proper experience as soon as possible.

The multiplexing equipment is normally unintelligent, i.e. it simply wraps up a number of cables into one single one. It is not impossible that intelligent multiplexers may appear which can act as simple circuit-switching equipment. This, however, remains to be verified.

The suitability of these systems for the transport of the current CERNET and INDEX systems is a vital question. The technical answer so far is that this is in theory possible. For CERNET in particular, a laboratory experiment is being carried out to verify this technical answer. In order to translate theory into practice it seems important to carry out a genuine implementation in a production environment, so as to experiment with the various parameters (cable type, speed, multiplexing).

As mentioned earlier, the 64 kilobit speed is a standard for digital telephone technology. It is an attractive possibility to include in the installed cable layout sufficient capacity for a telephone system. This has already been discussed with the CERN negotiator for PTT matters (A. Herz).

An alternative technology much in vogue elsewhere, especially in the United States, is that of Community Antenna (or Cable) Television. This is inherently a frequency division multiplexing system, normally using coaxial cable. It is a broadcast medium, in which the receiver tunes to the required frequency. Despite its apparent complete dissimilarity to G-700 series, it can be used for similar purposes, including digital transmission (although all of the transmission equipment relies on analogue techniques). It is likely that there will be a CATV coaxial cable in the tunnel, used for access control, raster scan waveforms and general information dissemination. The extent to which other types of service may use the cable is yet to be studied.

Transmission techniques for connections exterior to CERN are fairly limited and not likely to evolve rapidly. For leased lines to other institutes a maximum speed of 9600 bits/second is currently the limit. Links at 48 kilobits, such as that connecting CERN to the Swiss PTT network Telepac, are rather more expensive. One may hope in the future for digital lines running at 64 kilobits, but the time-scales are not clear. Satellite connections at high speed, such as that offered by Stella, will probably be too costly (in the short term at least) for such tasks as the transfer of experimental raw data.

There are other basic transmission technologies in use, especially for medium to high speeds. It is not possible to describe them all, and in any case they frequently belong in a closed domain such as a Local Area Network. However, some effort will be needed to keep up to date with them so as to be able to offer advice where necessary.

2. AVAILABLE PRODUCTS

It seems appropriate to give a short review of communications-related products commercially available now or in the near future. Relevant information includes technology, speeds, range, cost, support availability and so on.

2.1 X-25 products

Now that X-25 is a recognised standard a range of products is becoming available, some of which might be of interest. The CERNET X-25 gateway itself is based upon microprocessor-based hardware and software developed by industry (Borer Electronics) to CERN specifications. Such products should become widely available, thus lessening the need for CERN to either design and build in-house or collaborate with outside firms.

For terminal usage via X-25 there is a need for Packet Assembler/Disassemblers (PADs). These can be programs inside minicomputers acting as terminal concentrators, or programmed microprocessor boxes. CERN has already acquired a PAD from Gandalf for use in terminal access between CERN and the outside world, and have tested a similar product from Sesa.

The current CERNET X-25 program uses INDEX as an entry mechanism for access to the host computers on the site. As the various computer manufacturers develop their own X-25 software it may be preferable in certain cases to install this software. One advantage would be the multiplexing of many terminal connections onto a single physical line.

Equipment offered for X-25 users is normally directed towards use at speeds up to 9600 baud. For 48 kilobaud the requirement is rather more special; there are currently only two users having a 48 kilobaud connection to Telepac, of which CERN is one.

2.2 Ethernet

The most widely talked about LAN is the Ethernet. It is a CSMA/CD (Carrier Sense Multiple Access with Collision Detection) network. The basic medium is normally a special coaxial cable, although some trials on standard 50 ohm cable have been tried. The maximum distance (cable end to end) is 2.5 kilometers and the basic speed 10 megabits per second.

Ethernet hardware has been available since 1981 from several sources, such as Interlan, 3Com and Ungermann-Bass. European representation is normally available (e.g. Adcomp AG offer Interlan hardware, Sension Scientific in England offer their own brand, possibly OEM). DEC have recently announced Ethernet hardware availability.

The main costs of Ethernet are at present the transceiver box, attached to the cable either by cut and insert or as a tap, and the particular device controller in use, connected to the transceiver by a short transceiver cable. Typical current costs are:-

1. The transceiver. Around 1000 Swiss Francs, less in quantity, should drop appreciably as mass production begins.

2. A terminal interface controller. Sension offer a quad channel v24 interface (i.e. 4 terminals or other serial asynchronous devices at speeds up to 9600 baud) for 1000 pounds sterling.

3. A computer interface. The first such controllers were for the DEC Unibus and Q-bus. The US price from 3Com was 3000 dollars (small quantities) last year. Digital have announced their own Unibus controller (and also their own transceiver). The prices are supposed to be competitive.

It is certain that there will be a wide range of Ethernet hardware and software within the next few years. Other manufacturers, including CDC and Norsk Data, are taking Ethernet as a basis for their own networking and communications development.

2.3 Cambridge Ring

The Cambridge Ring, a slotted ring technique LAN, is being strongly supported by the UK. It will be marketed by several firms, including Logica, who give it the name of Polynet. The basic clocking rate of the ring is 10 megabits per second, although the techniques used limit the practical throughput by a factor of 2 or more. The theoretical circumference of the ring may be very large, but practical requirements tend to limit one to the same order of size as Ethernets.

The prices currently quoted by Logica, who are offering the Cambridge Ring under the name of 'Polynet' are rather high. A node costs 750 pounds sterling, a PDP DMA interface 3100 pounds sterling, a Multibus DMA interface 1350 pounds sterling, a program interrupt interface (PDP or Multibus) 350 pounds sterling and a monitor station 2500 pounds sterling. These prices must surely come down considerably.

2.4 Future LANs

No remarks on LANs could be complete without trying to crystal-gaze on the intentions of IBM. It is known that they have an interesting token passing ring working in their Zurich laboratory, as well as a very fascinating ultra-high performance block switch (like a packet switching node but very much more throughput). It is also said that IBM has paid millions of dollars for a licence for the token ring patent. There is likely to be an announcement fairly soon.

2.5 Metropolitan nets

This term, abbreviated to 'metronet', is meant to refer to the sort of network covering a metropolitan area. As such it matches the geographical size of LEP.

The only commercial offerings of such networks seem to be based on CATV techniques. Sytek, based in California, offer terminal access over ranges up to 50 kilometres using a CATV network. Network Systems corporation, who have marketed the Hyperchannel for some years, also envisage competing in the metronet field, probably using CATV technology to interlink Hyperchannels and Hyperbusses. Gandalf, the makers of INDEX, have plans to use CATV to extend their product line with a general system which they call Pluto.

The prospects of such networks, seen as a commercially available replacement for CERNET are still unclear, in that there is probably not a very large commercial demand for such high speeds over such long distances. Nevertheless it is a field which must be watched.

2.6 Personal Work Stations

The aspect of the personal work stations of particular importance for communications is the LAN to which they choose to attach, given that it is now unthinkable to consider them in a stand-alone environment. Of parallel importance is the way in which some kind of a gateway (hardware plus software) can be provided for such networks.

The Apollo comes with a dedicated LAN, called 'Domain', and software allowing the work station user to select as file base any other work station on the network. There are as yet no associated products to allow connection elsewhere, except possibly for the ubiquitous IBM Hasp Work Station emulator and the stated intention to provide an Ethernet connection. Therefore, the physical connection of a gateway Apollo, plus the necessary software, might need to be a CERN development project. More experience in the use of these devices is needed before the difficulty of such a development can be assessed.

Other personal work stations, such as the Perq and the Sun, come with an Ethernet interface. This should in principle simplify the task of interconnection, which could be either by programming one of the work stations or by having a separate gateway (a PDP running UNIX?). Again, there are not many details on commercial offerings, but something must surely come.

3. COMPLETE SYSTEMS

By 'complete systems' is meant systems which offer a complete user function, rather than those which simply provide the tools with which to create the function. The definition is not particularly precise; for those familiar with the ISO Open Systems Interconnection model it may approximate to levels 5 and above, and on occasions to level 4.

Not surprisingly, most complete systems are built around the products of a single manufacturer. Classical examples are the Systems Network Architecture (SNA) of IBM and the Digital Network Architecture (DNA) of DEC. Other minicomputer manufacturers offer similar products.

All of these manufacturers are moving towards interconnection via X-25 links or networks for long-distance connections. They tend to have their own protocols, plus the ability (for non-IBM) to simulate some IBM protocols. Thus, for instance, DEC is basing the connection strategy on a combination of X-25 and Ethernet.

It can thus be seen that the manufacturers are climbing the rungs of the OSI model, and have got to level 3. For the step to level 4 there is a hope that the ECMA transport protocol may be adopted between now and around 1985. The step to higher levels is very much more open to question.

3.1 General Systems

Some complete systems are the type that may be installed on any of the range of computers for which they are designed. That is to say that one may have DECNET, including file transfer, job enquiry and so on, on any VAX (or perhaps PDP/11) running a standard operating system. Computer mail systems, and even more sophisticated office automation systems such as Digital's Office Plus or IBM's Profs, could also be widely available. The problem then arises of the interconnection of these systems over CERN.

3.2 Specialised Systems

There are some kinds of complete system which are not normally run as an add-on service, but rather as a specialised service running in one particular location. The current projects for the LEP Data Base and Computer-Aided Engineering are examples of such systems. Other products, such as teleconferencing, project management systems, are or will be available. Here the problem is to provide for the access mechanism allowing the user (i.e. a display and a keyboard attached to something or other) access to the required service.

3.3 Problems

The essential point in the acquisition of a 'complete system' is that experience shows that there is a high probability of requiring interconnection between the complete system and other systems. The projects for LEP data base and CAE were bought with this in mind; the various specialist products for office automation within divisions (Wang, Philips, Siemens) were not.

APPENDIX 5

W.G.2 Study Group 1 Interim Report
Suitability of CCITT G-700 TDM Systems for CERN

Author: P. Wolstenholme
Date: 6 Dec 82
Version: 8
Status: Final

1. QUALITATIVE FACTORS

1.1 TRANSMISSION SYSTEMS

A variety of transmission services is required at CERN, at line speeds in the range 300 to at least 3,000,000 bits/second. Over the rather long transmission distances of LEP - up to about 12 km for the longest links - there is certainly a need to use equipment which has been developed by experts in line transmission, and the most obvious candidate systems are the digital transmission systems used in modern telecommunications systems, to the G.700 series of C.C.I.T.T. recommendations [1].

Transmission systems operating at 2, 8.5, 34 and 140 Mbits/second are available from many firms in the CERN member states. They are all designed to comply with the G.703 interface recommendations, and are thus compatible and interchangeable. For several transmission systems over copper pairs the transmission distances, between regenerators, are well matched to the distance between LEP access shafts of 3.4 km

Potential advantages in using these transmission systems include the following:

- competitive tendering within CERN member states
- availability of modern but compatible equipment assured for the next 30 years at least.
- standard test equipment is available

- all these systems have a measure of redundancy which permits error rates to be monitored irrespective of the transmitted data and framing employed.
- although CERN development is not forbidden we do not need to do much, and can profit from the experience of others. (The situation here is comparable to that of the CAMAC standard.)

There are, naturally, a few disadvantages. In the LEP tunnel we shall not be able to use optical fibre transmission, and cables will be required. Twisted pair cables can be employed at 2 and at 8.5 Mbit/s speeds, but coaxial cables are required for 34 Mbit/s and higher, and preferred at 8.5 Mbit/s. These cables are required to have very low near-end cross-talk (i.e. low interference to reception from the adjacent transmission unit). Cables with very good screening, i.e. solid outer tubes rather than braid, and with a low-loss helical membrane or similar support for the inner conductor are specified, and these are expensive.

The entire system for multiplexing and transmission requires that the signals being transmitted conform to special requirements, and that the clocks are either very precise indeed or come from the multiplexing equipment itself.

1.2 MULTIPLEXERS

Multiplexing equipment is used to permit high-speed channels to be employed to transmit many channels at low speeds - typically the 64 kbit/s of the speech channel. If we need to support a variety of services at speeds from a few kbit/s to about 8 Mbits/s, then the multiplexing of these services into smaller numbers of high speed channels will usually give some advantages in cost, maintainability and future flexibility of the communications system. This is not always true, as we have to pay for the multiplexing equipment.

Advantages of using standard digital multiplex equipment, apart from the possibility of using the transmission equipment mentioned above, are similar to those of using the transmission equipment. They include multiple European sourcing, long term viability, availability of standard test equipment, etc. Traffic at rates above 2 Mbit/s can be multiplexed easily with currently available standard equipment, but the situation regarding lower speed traffic leaves much to be desired.

The major difficulty in using the above systems at the lower speeds is that most of the equipment designed to date has been designed around the telephone speech channel. Very little commercial equipment of general purpose nature is on the market today. This situation will gradually improve over the next few years, and we also have a pilot project [7] at ISR to develop some suitable equipment which could eventually be commercialised.

Although the primary multiplexers were developed for transmission of digitised speech only, at 64 kbit/s, it is possible to adapt or design multiplexing equipment to transmit data at a mixture of data rates, preferably small or binary multiples of 64 kbit/s.

1.3 INTERFACES

Until the emergence of the new universal ISDN [4] interface, not too likely in the LEP construction time scale, the IS 4903 interface [5] is recommended [2] by the CCITT for line speeds in the 20 kbit/s to 10 Mbit/s range, specifically for X.25 data networks. It employs X.24 and X.27 (RS422) electrical standards. There are proposals by various telephone administrations to use it in a preliminary digital subscriber service, at 64 kbit/s.

If CERN adopts such standards internally, on as wide a scale as practicable, it will not be difficult to take advantage of digital services as they become available for our links to other laboratories around the world.

1.4 SWITCHING

Any system which is multiplexing low speed channels on to a high speed carrier is easily extended to provide a switching function, by changing the channel allocations in the multiplexers. This can be done infrequently, in which case it serves as a reconfiguration technique for maintenance, adding new users, etc. If done on a more dynamic basis, it becomes possible to allow switching at the request of the user, as in the INDEX system. This is achieved by connecting channels at one side of the system to individual users, multiplexing them in some way, and connecting the other side, after demultiplexing, to the various services. By using the appropriate techniques to control the allocations of channels at the two sides of the system the required services are connected to users, on a one to one basis.

The use of time division multiplex facilitates the switching system implementation, and it is not expensive. There are even integrated circuits which can connect 240 subscriber 64 kbit/s channels (on 8 2 Mbit/s lines) to 240 outgoing channels (again on 8 2 Mbit/s lines) in any required fashion, under the control of an address table loaded by a microprocessor, and these circuits will also function in cases of channels using more than one 64 kbit/s speech path. This technology is being introduced into the current generation of electronic telephone exchanges (see section 4 below).

The essential feature of such switching systems is their transparent nature. When a connection has been set up, there is a direct path from one input port to one output. (In fact multiple outputs are easily set up, but the converse of merging multiple inputs into one output is difficult.) The transparent nature of the connection may be valuable in some circumstances, but there may be requirements which are not easy to satisfy. These might include multiple logical connections, as in a packet-switched data network.

1.5 IMPLICATIONS OF SYNCHRONOUS INTERFACES

An important feature of synchronous multiplexing and transmission is that it avoids some of the problems encountered in INDEX. Once the data has been put into synchronous format, it can be transmitted through very many stages of multiplexing and switching with no cumulative introduction of jitter, and therefore a low risk of errors in transmission from this cause.

The synchronism can be achieved in an adaptor, if we need to use an asynchronous terminal of the type common at CERN. The adaptor might be a "Triple-X PAD" (Packet Assembly-Disassembly unit) [6] if interfaced to a modern local area network. The best way is, of course, to avoid using the asynchronous interface in other than purely local terminal-to-computer configurations.

2. ECONOMICS

The economics of transmission systems at each level are not too controversial. The equipment is constructed to high standards of reliability and maintainability, for which one has to pay, but is not specially developed for CERN and is not specially expensive. Prices range from about 5000 SF for each end of an 8.5 Mbit/s duplex transmission system to about 20000 SF per station at 34 Mbit/s for an optical fibre link. In the CERN context these can not represent a very large sum of money, but the crucial issue is the selection of the appropriate transmission speeds for the physical links, bearing in mind the costs of cables and their installation on the one hand and the cost of multiplexing equipment on the other.

It should be possible to transmit over distances of up to about 5 km at low speeds up to 2 Mbit/s using rather cheap cable costing about 300 to 500 SF per pair per km. At higher speeds the twisted pair cables are difficult to use, and 8.5 Mbits/s over 3.4 km is a practical limit for rather high quality cable.

Coaxial cable to telecommunications specifications [3] is expensive. A 4-core coaxial cable to G.622 (1.2/4.4 mm) will cost about SF 10 per metre, i.e. 2500 SF per core per km. Such a cable will support 34 Mbit/s over 4 km spans between regenerators, or 140 Mbit/s over 2.05 km. It could be installed in the LEP tunnel, initially without regenerators, and operated at 34 Mbit/s. Regenerators installed in the alcoves, protected from radiation damage, could be installed either to upgrade the capacity or to improve margins to cure an eventual interference problem. The next larger cable, again with 4 pairs, is to recommendation G.623 (2.6/9.5 mm) and would cost about SF 25 per metre. This can support 140 Mbit/s over 4.6 km spans.

Fibre optic cable costs about 2 SFr. per meter per fibre, and is becoming cheaper each year. Using semiconductor laser sources, it can support transmission at up to 500 Mbit/s over a few km. or alternatively 34 Mbit/s over the LEP site using LED transmitters, which should be more reliable than lasers.

A factor not to be neglected is the cost of a suitable amount of spare equipment and cabling, if the availability of the service is to be high and links on hot standby are to be required.

A multiplexer (duplex) connecting four channels at 2 Mbit/s to one at 8.5 Mbit/s costs about 7000 SF while the next level up, connecting 4 channels at 8.5 Mbit/s to one at 34 Mbit/s costs about 12000 SF. Low speed multiplexers are more expensive as they are usually dealing with many channels. For example, a multiplexer for 30 channels at 64 kbit/s (interfacing to a 2 Mbit/s G.732 system) will cost about 25 to 30 kSF, while one dealing with a smaller number of channels at various speeds and standards might cost just under 20 kSF.

On the CERN site the transmission distances are short, and the benefits of multiplexing can easily be outweighed by the cost. Taking a 2 Mbit/s CERNET link as an example, a direct 12 km link over a single quad cable might cost about the same as the use of part of an 8.5 Mbit/s circuit over a more expensive cable, while the use of part of a 34 Mbit/s optical fibre link will in fact be cheaper over this distance. The choice will be determined in practice by factors such as ease of expansion, need to provide a variety of services in a homogeneous framework, maintenance policy etc.

3. APPLICATIONS IN CERN

3.1 HIGHER SPEEDS

There are several potential applications for high speed digital transmissions in CERN, particularly associated with the LEP project. Data links for the control system and for the experiments are required, at 2 or 8.5 Mbit/s. Other classes of transmission such as closed-circuit television may also be considered. High speed links may well be used for low speed traffic, via suitable multiplexers. The economics of transmission right across the LEP site will no doubt favour the use of fibre optic systems at 34 Mbit/s, possibly upgraded to 140 Mbit/s when the traffic has grown sufficiently to fill the initial links. In the tunnel there is a difficult choice to be made, between coaxial and twisted pair cables, although it is clear that coaxial cables will offer a much greater potential capacity, even up to 140 Mbit/s using regenerators in the alcoves.

3.2 LOW SPEED CHANNELS

Requirements for low speed channels range from the present terminal traffic at 4800 baud to the support of intelligent work stations, microprocessor development systems etc. for which higher speeds, in the range 48 kbit/s to about 300 kbit/s, will be required in the near future. These channels could all be multiplexed on to 2 Mbit/s transmission channels, and further multiplexed whenever that would be convenient and cost-effective.

It must be emphasised that the equipment for multiplexing and transmission is essentially transparent in nature, and replaces lengths of wire. It deals neither with connections, routing nor any other communications protocols: those are at higher levels of the ISO model. Accordingly, it is useful and economic only over fairly long distances.

It will certainly be feasible to use T.D.M. links for terminal similar traffic. Multiplexers with the required characteristics are starting to appear on the market. It must be reiterated that such equipment deals at present with multiplexing point to point links over fixed routes, and even with the addition of a switching function it does not solve all the problems of future CERN services. If future development of CERN services leads to the quasi-permanent connection of the typical user to two or more hosts, e.g an IBM service, the Oracle data base and a message system, a dynamic routing will be needed and only a unified CERN network strategy incorporating packet-switching for a range of services will satisfy such a requirement.

The software switching and multiplexing provided by packet-switched networks complements the hardware features provided by TDM systems: both are needed. A CERN-wide network providing many services might need to be made up of a two level physical network, in which a high speed network would be used to interconnect a number of local area networks. This is not directly relevant to the work of Sub-Group 1, but is mentioned on account of its influence on the multiplexing strategy. Only the high speed network will require long distance transmission, and there will be little need to multiplex low speed channels.

3.3 DIGITAL TELEPHONE EXCHANGES AT CERN

It is often suggested that CERN might one day install a modern telephone exchange providing full digital interconnection at 64 kbit/s (or a higher rate) from every place where we now have a telephone instrument. Although it might be possible to base part of our communications strategy on use of such a telephone system, there would be two major problems with such an approach. It is not likely that the required locations of digital devices (terminals, printers, specialised displays etc.) will map exactly onto the telephone requirement one for one. Furthermore, the telephone system only provides for switching and transmission on a point-to-point basis, rather than the dynamic and multiple logical connections really needed in the long term.

The transmission facilities offered by a digital telephone system could very well be used to interconnect low cost local-area networks, or isolated devices, to the main CERN networks in a flexible way. This should be done at either 4800 baud (INDEX compatible) or 64 kbit/s as appropriate, and a dial-up service would seem to present the most useful possibility. Some fixed point-to-point links might be permitted via the telephone system, but a standard and independent CERN-wide communications facility would be recommended for permanent links.

A typical example of the occasional user best connected via the telephone system might be the isolated secretary with a word processor. Even this example would be inappropriate were the secretary to be connected to her Group via a local area network having a message facility.

Use of a CERN internal telephone system as described above will readily be extended to use the international digital network when it arrives, and allow our leased connections, upgraded to 64 kbit/s working, to be as easy to use as the local lines.

4. CONCLUSIONS

4.1 PRESENT SERVICES

It is feasible to run CERNET at 2 and perhaps 8.5 Mbit/s over standard T.D.M. channels via the G.703 series of interfaces, and practical tests are under way at the present time to support this statement.

INDEX and similar traffic can employ T.D.M. transmission channels using presently available equipment. The equipment would have to be very carefully selected, as much of the commercial multiplexing equipment now on the market relies on the "statistical multiplexer" principle. This is unsuitable for CERN terminal traffic, as our main computers have no flow control protocol.

T.D.M. systems are only economic over large distances such as those encountered in the LEP site, but an eventual integrated CERN-wide system might apply them to shorter distances such as between the computer centre and the North Area. Consideration should be given to installation of a coaxial cable on a route from Building 513 to the North Area: a distance of about 3.8 km. Such an installation would be useful in the short term, running at 8 Mbit/s for CERNET and other services, and could be tested at 34 Mbit/s. The route, with a possibility of installing regenerators at 1.3 km intervals, closely corresponds to the situation of a cable around one octant of the LEP main ring, and a valuable field trial would thus be made. As this route will take most of the traffic between the present computer centre and the LEP experiments, the installation would have long term utility.

Optical fibre transmission systems are already cheaper than coaxial cable systems, and will allow us to accommodate future increases in data rates in an economical way. The main links between the Meyrin and Prevessin sites may ultimately be based on fibre optics, and it could be useful to install a cable on this route as soon as the duct becomes available.

4.2 FUTURE SERVICES

The Study Group has not yet discussed possible future services in any detail. Another important matter regarding services such as television distribution is the choice between digital line systems using TDM and the alternative analog (FDM) wide band systems currently in use for CATV and other purposes. This has not been studied, but should receive some attention in order that the best overall choice can be made for CERN The TDM system is at its most economic over relatively long transmission distances, and traffic in the rather small Meyrin and Prevessin sites might be efficiently handled on a broadband FDM communications system. Studies ought to be made of various possibilities so as to use either one system or a combination of both systems.

5. REFERENCES

1. The International Telegraph and Telephone Consultative Committee (CCITT) VIIth. Plenary Assembly (Geneva) November 1980. Published as the "CCITT Yellow Book" by I.T.U. - Geneva, 1981. Volume III Fascicle III.3 Transmission Systems and Multiplexing Equipment.

2. ibid. Volume VIII Fascicle VIII.2 Data Communication Networks... Terminal Equipment and Interfaces.

3. ibid. Volume III Fascicle III.2 Transmission Media.

4. "Pathways to the Information Society" Proceedings of the Sixth International Conference on Computer Communication, London, September 1982. North Holland Publishing Company. Several papers, pages 13 to 60 inclusive, dealing with the Integrated Services Digital Network (ISDN).

5. "Data Communication - High level data link control procedures - Frame structure" International Standard 3309 (1976) International Organisation for Standardisation, Geneva, Switzerland.

6. F. Fluckiger: Overall Presentation of the CERN X.25 Program PAP/05/03/Nov.02 1981.

7. P. Wolstenholme and A. J. Vezis: Digital Communications Systems for LEP. LEP Controls Note 17, November 20, 1981.

APPENDIX 6

Summary of WG2 Subgroup 2 Results
(Characterisation of protocol services)

Author: B.Carpenter, D.Samyn, for SG2
Date: 28 Oct 82
Version: 4
Status: Final

1. INTRODUCTION

This report attempts to summarise the results of the investigations of WG2 Subgroup 2 (Protocol Services). A major source for the information used in this summary is the body of documents assembled by WG2, for which no individual references are given. In addition, information gained from the questionnaire sent to several LEP collaborations has been incorporated as appropriate.

The report is presented as follows:

1. identification of domains of application

2. identification of individual protocol services

3. characterisation of each service for each domain

The domains have some correspondence to the 'users' identified by Subgroup 3; the services have some correspondence to the levels below 7 of the ISO/OSI Reference Model; but these correspondences are by no means precise. The intention is instead to ensure that all users and high-level services identified by SG3 are catered for by the domains and protocol services covered here.

To shorten the presentation of the results, each domain, service and characterisation is denoted by a mnemonic when it is defined.

2. DOMAINS OF APPLICATION

Notes:

(a) "Controls" includes accelerator controls and SB services.

(b) "Centralised computers" refers to mainframes and other computer-centre facilities (CAD/CAM, database, and so on). "Decentralised computers" refers to facilities outside the computer centre providing general or specialised services not mentioned elsewhere (such as ADP or library services).

(c) "Office automation" is interpreted to include mail and document preparation, not only secretarial facilities, and possibly some personal computing as well. Note that communication between secretarial staff and technical staff (the latter probably not using office automation terminals) is probably more important than that between different secretaries.

(d) Personal workstations are not included as a domain; allowance is made for them among the services.

1. C-NETW: Controls: real-time traffic in computer network
2. C-INTF: Controls: traffic to interface (may be dumb or smart)
3. C-CERN: Controls: traffic to rest of CERN
4. E-HARD: Experiments: event data handling dedicated hardware
5. E-LOCL: Experiments: local minicomputer control/monitoring/supervisory activities
6. E-MPUS: Experiments: microprocessor interconnections
7. E-CERN: Experiments: traffic to rest of CERN
8. E-OUTS: Experiments: traffic to outside CERN
9. S-CENT: Software development: for centralised computers
10. S-DECE: Software development: for decentralised computers
11. S-XCOM: Software development: for cross-supported computers
12. S-XMPU: Software development: for cross-supported microprocessors
13. G-CENT: General computing: access to centralised computers
14. G-DECE: General computing: access to decentralised computers

15. G-CERN: General computing: traffic between various centralised and decentralised computers

16. G-OUTS: General computing: traffic to outside CERN

17. O-LOCL: Office automation: local traffic (within division)

18. O-CERN: Office automation: inter-division traffic and traffic within a divsion between secretarial and technical staff.

19. O-OUTS: Office automation: traffic to outside CERN

3. PROTOCOL SERVICES

Notes:

(a) No comment is made about implementation, especially as to whether the services are provided by circuit switching or by packet switching (either datagram or virtual circuit protocol).

(b) No comment is made about gateways; the need for gateways can in some cases be deduced from later chapters; in other cases, from the SG3 models.

(c) Mail services (although very important) are not considered to constitute a protocol service for the purposes of this report.

1. T-T: task/task communication

2. TER: (remote) terminal

3. FAP: remote file access protocol

4. DLL: down-line load/ up-line dump

5. RJC: remote job control

6. RPC: remote procedure call

7. PAG: paging (virtual memory)

4. CHARACTERISATION OF SERVICES IN EACH DOMAIN

This chapter attempts to list various characteristics of each service in each domain, as follows:

1. Service needed (YES/NO)
2. N: Number of stations (physical connections, ports, subscribers)
3. ADR: Average data rate per station (bytes/second)
4. PDR: Peak data rate (bytes/second over given interval)
5. MAX: Maximum distance (metres)
6. MTBF: Desired MTBF per station (days)
7. MTTR: Desired MTTR (hours)
8. NAW: Night and weekend service essential (YES/NO)
9. (Comments)

5. RESULTS

We have attempted to evaluate each of the above characteristics for each service in each domain. Obviously the evaluations are estimates in many cases, although in general they are based on WG2 documents or on the questionnaire. The detailed (but inevitably incomplete and approximate) table of results is appended to this note. In general, the data rates requested by the experiments (but not the other characteristics) in our opinion tend to err on the high side. Also, the estimates given are for late in the LEP construction period and beyond. At least linear growth from the current situation can be expected up to that time.

The following sections present qualitative results from the questionnaire (replied to by ALEPH, DELPHI, L3 and OPAL).

5.1 SIZE OF EXPERIMENTAL COMPUTING FACILITIES

Each experiment will have up to 30 terminals, 15 graphics terminals, 10 intelligent workstations, 10 minicomputers, 50 emulators and 200 microprocessors.

5.2 WORKSTATIONS (TERMINALS ETC.)

- There will be a large amount of alphanumeric terminals (15 to 30) at the experimental area;

- The graphic terminal needs vary widely between experiments. - The use of intelligent workstations (PERQ, APOLLO) depends on the price.

In principle, a switchable (INDEX like) workstation system is desirable only if such a system is cheap, and offers greater security and greater speed than dedicated workstations. In any case, some workstations must be dedicated (1) during the development stage, (2) on some critical subsystems of the experiment.

5.3 ACCESS TO PERIPHERALS

A uniform system for access to shared peripherals (disc, MT, printers, etc) is only acceptable (1) if the manufacturers provide HW and SW for connecting peripherals on any CPU, (2) or, in case of network access, if speed, security and reliability is sufficient and if transparent file handling is provided.

There are some doubts whether this will be the case except for low speed peripherals (mainly printers). Furthermore, during the development stage, the computers will be distributed around the CERN site or in the home labs and will need their own (cheap) discs and printers.

Backup problem: people plan to backup their discs locally using the manufacturers' backup procedures or on a local data base. Some have experience of the current CERNET backup procedure which is considered to be too slow. If a central backup facility is provided, the speed should be higher and it should be able to handle 50 to 100Mbytes per mini and per week. This is equivalent to a total of several gigabytes per week for all experiments together. One backup operation should take about 15 minutes.

5.4 MICRO-PROCESSORS

The use of micros will be widespread: 50 to several hundred 68000s or 6809s (20KB to 512KB memory size each). They will perform various functions, in addition to event data transfer: debugging, collecting event statistics, routine performance checking, parameter controlling, up/down-line loading, command/messages exchanging with minis, etc. The size of the event data exchanged will vary from 10 to 300 Kbyte. A speed of 100K to above 1M byte/sec is required for critical data communications.

A single all-purpose inter-micro network is not considered as a plausible solution. Rather, a dual solution is considered by the various experiments:

- cheap LAN (or RS232C connections) for low speed communications: debugging, up/down-line loading, command/message exchange. Certain experiments requested very high speed (300 Kbyte/second) on the 'cheap' LAN: this is technically unrealistic today.

- hardware such as CAMAC, FASTBUS, shared memory buffers for high speed transfers.

5.5 ACCESS TO THE EXTERNAL WORLD

This concerns the rest of CERN, outside CERN, and the LEP controls system.

The number of 'outside' users of the experiments' computing power will be very high:

- 10 to 50 users in CERN per experiment,

- 15 to 50 users in Europe per experiment,

- 2 to 30 users in North America per experiment,

- 2 to 20 users elsewhere in the world per experiment.

Therefore, all kind of services are required: alphanumeric and graphic terminal access, remote printing, remote file access, job submission, real time data transfer (remote test of equipment), computer aided design, project management aids, software management (distributed software), documentation management, data base access and mail (most crucial!).

All experiments request substantial efforts to be put into communications with home institutes.

LEP control: a bi-directional on-line connection to the LEP machine control system is required for each experimental area and its use will be frequent. The requirements (data size, type and speed) are for the moment unclear, but will be relatively modest.

5.6 EXPERIMENTAL LAN

A unifying idea in the minds of many people is that of a LAN of the Ethernet class for each experiment, connecting the minis, interfaced to terminals via PADs and with gateways to the CERN backbone (hence to outside CERN). It is considered unlikely, for economic reasons, that this LAN can extend to the micros: hence either the cheap LAN mentioned above, or RS232 lines via PADs, will be needed for micros.

6. IMPLICATIONS FOR CERN COMMUNICATION NEEDS

On current thinking, LEP experiments have serious implications for CERN future communication needs:

1. Experiments will require extensive and comprehensive data connections to the rest of CERN and to many outside institutes (not only in Europe and North America; Russia, China and Japan have all been mentioned). Remote interactive graphics have been seriously requested by several experiments, although the feeling of most technical experts is that this is not the best approach; and long-haul protocols for graphics do not exist.

2. External connections, via X25, are of crucial importance and must be available all over the CERN/LEP sites. At least 100 simultaneously established virtual circuits must be envisaged. Higher-level application services, especially mail and software development tools, must work well over these connections, and logically should be centralised for and at CERN;

3. Each experiment will require a LAN of the Ethernet class for a few tens of stations (minicomputers, PADs, gateways, peripheral servers);

4. Each experiment will require a cheap LAN (or equivalent) for microcomputers, with up to a few hundred stations, 200 metres. Speed as high as possible! Can only be combined with the preceding LAN if cheap enough.

5. The experiments tend to see CERN's role as providing the X25 service, possibly the two levels of LAN, and integrated protocol and high-level services across these systems and their various gateways. This presumably includes operational support as well as initial selection and implementation.

6. Transfer of event data through the various layers of an experiment, up to the data acquisition computer(s), and data for on-line processing, is not to be treated as a communication problem for WG2.

7. Experiments will require two-way connections to the LEP machine control system.

8. Electronic , mail and related facilities are crucial to LEP collaborations (and must be available internationally).

9. Experiments hope for a high level of support from CERN (recommended and fully supported standards; out-of-hours service; etc)

10. The range of communications protocol services needed for controls and for experiments are not particularly different in any major characteristic. However, LEP and LPI (Lep Pre-Injector) controls both have significantly nearer deadlines than the LEP experiments and so

can probably not wait for a CERN or industry standard local network to emerge.

7. CERN-WIDE PROTOCOL NEEDS

In general most protocol services are needed in many domains, in particular TER (remote terminal) and FAP (file access protocol) which are needed essentially everywhere. DLL (down-line load) and RJC (remote job control) are needed in about half of the domains. At first sight T-T (task-task communication) is less widespread, but this is an illusion as it will in practice be the basis for most other services. Only two services are truly restricted: PAG (paging) for certain types of proprietary office-automation and personal-computing LANs; and RPC (remote procedure call) for controls.

The fact that the services have different mundane requirements (number of stations, data rate, distance and reliability) in different domains may affect choice of hardware and driver-level software; but otherwise there seems to be no good technical reason why a common set of protocols cannot be adopted as CERN standards. These can initially be regarded as gateway protocols for use at the LAN/LAN or LAN/backbone interface; later they could become implementation protocols for new LANs; or they might be superseded by newly adopted protocols in a controlled way.

It is not the business of SG2 to attempt any protocol definitions, but we do recommend that

1. It is decided to adopt standards for T-T, TER, FAP, DLL and RJC protocols (assuming RPC and PAG remain private). A mail service should be defined in addition.

2. That these protocols be independent of hardware and driver levels.

3. That these definitions should be available by middle 1983.

4. That a mechanism be established to test and validate implementations of these standards.

Experiment	Contacts
ALEPH	W. von Rueden
DELPHI	R. Dobinson
L3	H. Newman (CALTECH)
OPAL	N. Gee (RAL)

			N #STATIONS	ADR BYTES/S	PDR BYTES/S	MAX METRES	MTBF HOURS	MTTR HOURS	NAW SERVICE	COMMENTS
C-NETW	T-T	Y	100	50-2000	10K	10K	5K	1	Y	ROBUST TOPOLOGY, PRIORITY SCHEME
	TER	Y	100	10	1K	10K	1000	1	N	
	FAP	Y	*	*	*	*	*	*	*	
	DLL	Y	*	*	*	*	*	*	*	
	RJC	N	-----	-----	-----	-----	-----	-----	-----	
	RPC	Y	*	*	*	*	*	*	*	
	PAG	N	-----	-----	-----	-----	-----	-----	-----	
C-INTF	T-T	Y	100	1-2000	10K	1K	100K	1	Y	
	TER	Y	10	1	100	1K	1000	1	N	
	FAP	Y	*	*	*	*	*	*	*	
	DLL	Y	*	*	*	*	*	*	*	
	RJC	N	-----	-----	-----	-----	-----	-----	-----	
	RPC	Y	*	*	*	*	*	*	*	
	PAG	N	-----	-----	-----	-----	-----	-----	-----	
C-CERN	T-T	Y	3	10	10K	20K	1000	2	Y	
	TER	Y	10	10	1K	20K	1000	1	N	
	FAP	Y	*	*	*	*	*	*	*	
	DLL	N	-----	-----	-----	-----	-----	-----	-----	
	RJC	Y	*	*	*	*	*	*	*	
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	N	-----	-----	-----	-----	-----	-----	-----	

* MEANS 'INCLUDED IN OTHER ESTIMATES'
----- " 'NOT APPLICABLE'
? " 'DOUBTFUL - NEEDS FURTHER STUDY'
BLANK " 'UNKNOWN'

C-NETW: CONTROLS: REAL-TIME TRAFFIC IN COMPUTER NETWORK
C-INTF: CONTROLS: TRAFFIC TO INTERFACE (MAY BE DUMB OR SMART)
C-CERN: CONTROLS: TRAFFIC TO REST OF CERN

T-T: TASK/TASK COMMUNICATION
TER: (REMOTE) TERMINAL
FAP: REMOTE FILE ACCESS PROTOCOL
DLL: DOWN LINE LOAD/ UP LINE DUMP
RJC: REMOTE JOB CONTROL
RPC: REMOTE PROCEDURE CALL
PAG: PAGING (VIRTUAL MEMORY)

Y/N: SERVICE NEEDED
N: NUMBER OF STATIONS
ADR: AVERAGE DATA RATE PER STATION (BYTES/SECOND)
PDR: PEAK DATA RATE (BYTES/SECOND)
MAX: MAXIMUM DISTANCE
MTBF: DESIRED MTBF PER STATION (HOURS)
MTTR: DESIRED MTTR (HOURS)
NAW: NIGHT OR WEEKEND SERVICE ESSENTIAL (Y/N)

TABLE 1 : CONTROLS

			N #STATIONS	ADR BYTES/S	PDR BYTES/S	MAX METRES	MTBF HOURS	MTTR HOURS	NAW SERVICE	COMMENTS
E-HARD	T-T	Y								DEDICATED HARD.
	TER	N	-----	-----	-----	-----	-----	-----	-----	
	FAP	N	-----	-----	-----	-----	-----	-----	-----	
	DLL	N	-----	-----	-----	-----	-----	-----	-----	
	RJC	N	-----	-----	-----	-----	-----	-----	-----	
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	N	-----	-----	-----	-----	-----	-----	-----	
E-LOCL	T-T	Y	50	?	2M?	200	5K	1	Y	
	TER	Y	50	10	1K	200	1000	0.5	N	
	FAP	Y	*	*	*	*	*	*	*	
	DLL	N	-----	-----	-----	-----	-----	-----	-----	
	RJC	Y	*	*	*	*	*	*	*	
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	N	-----	-----	-----	-----	-----	-----	-----	
E-MPUS	T-T	Y	200	?	300K?	200	100K	1	Y	
	TER	Y	10	10	1K	200	1000	1	N	
	FAP	Y	*	*	*	*	*	*	*	
	DLL	Y	*	*	*	*	*	*	*	
	RJC	N	-----	-----	-----	-----	-----	-----	-----	
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	N	-----	-----	-----	-----	-----	-----	-----	
E-CERN	T-T	Y	3	?	1M?	20K	1000	1	Y	REQUESTED BACKUPS:
	TER	Y	50	10	1K	20K				1GBYTES/15MINS
	FAP	Y	*	*	*	*	*	*	*	FIGURES PER EXPT.
	DLL	N	-----	-----	-----	-----	-----	-----	-----	S=4*NO. OF STATIONS
	RJC	Y	*	*	*	*	*	*	*	
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	N	-----	-----	-----	-----	-----	-----	-----	
E-OUTS	T-T	N	-----	-----	-----	-----	-----	-----	-----	
	TER	Y	50	10	1K	5000K	1000	4	N	
	FAP	Y	20	30?	1K	5000K				
	DLL	N	-----	-----	-----	-----	-----	-----	-----	
	RJC	Y	*	*	*	*	*	*	*	
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	N	-----	-----	-----	-----	-----	-----	-----	

* MEANS 'INCLUDED IN OTHER ESTIMATES'
----- " 'NOT APPLICABLE'
? " 'DOUBTFUL - NEEDS FURTHER STUDY'
BLANK " 'UNKNOWN'

E-HARD: EXPERIMENTS: EVENT DATA HANDLING DEDICATED HARDWARE
E-LOCL: EXPERIMENTS: LOCAL MINICOMPUTER CONTROL/MONITORING/SUPERVISORY
E-MPUS: EXPERIMENTS: MICROPROCESSOR INTERCONNECTIONS
E-CERN: EXPERIMENTS: TRAFFIC TO REST OF CERN
E-OUTS: EXPERIMENTS: TRAFFIC TO OUTSIDE CERN

T-T: TASK/TASK COMMUNICATION
TER: (REMOTE) TERMINAL
FAP: REMOTE FILE ACCESS PROTOCOL
DLL: DOWN LINE LOAD/ UP LINE DUMP
RJC: REMOTE JOB CONTROL
RPC: REMOTE PROCEDURE CALL
PAG: PAGING (VIRTUAL MEMORY)

Y/N: SERVICE NEEDED
N: NUMBER OF STATIONS
ADR: AVERAGE DATA RATE PER STATION (BYTES/SECOND)
PDR: PEAK DATA RATE (BYTES/SECOND)
MAX: MAXIMUM DISTANCE
MTBF: DESIRED MTBF PER STATION (HOURS)
MTTR: DESIRED MTTR (HOURS)
NAW: NIGHT OR WEEKEND SERVICE ESSENTIAL (Y/N)

TABLE 2 : EXPERIMENTS

			N #STATIONS	ADR BYTES/S	PDR BYTES/S	MAX METRES	MTBF HOURS	MTTR HOURS	NAW SERVICE	COMMENTS
S-CENT	T-T	N	-----	-----	-----	-----	-----	-----	-----	
	TER	Y	500?		10K	20K			N	
	FAP	Y	*	*	*	*	*	*	*	
	DLL	N	-----	-----	-----	-----	-----	-----	-----	
	RJC	Y	*	*	*	*	*	*	*	
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	N	-----	-----	-----	-----	-----	-----	-----	
S-DECE	T-T	N	-----	-----	-----	-----	-----	-----	-----	
	TER	Y	500?		10K	20K			N	
	FAP	Y	*	*	*	*	*	*	*	
	DLL	N	-----	-----	-----	-----	-----	-----	-----	
	RJC	N	-----	-----	-----	-----	-----	-----	-----	
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	?			10M?					PAGING FOR WORKSTATIONS?
S-XCOM	T-T	N	100?		50K	20K			N	
	TER	Y	*	*	*	*	*	*	*	
	FAP	Y	*	*	*	*	*	*	*	
	DLL	Y	*	*	*	*	*	*	*	
	RJC	Y	*	*	*	*	*	*	*	
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	?								
S-XMPU	T-T	N	100?		50K	20K			N	
	TER	Y	*	*	*	*	*	*	*	
	FAP	Y	*	*	*	*	*	*	*	
	DLL	Y	*	*	*	*	*	*	*	
	RJC	Y	*	*	*	*	*	*	*	
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	?			10M?					PAGING FOR WORKSTATIONS?

```
  *      MEANS   'INCLUDED IN OTHER ESTIMATES'
-----      "     'NOT APPLICABLE'
  ?        "     'DOUBTFUL - NEEDS FURTHER STUDY'
BLANK      "     'UNKNOWN'

S-CENT: SOFT. DEV.: FOR CENTRALISED COMPUTERS
S-DECE: SOFT. DEV.: FOR DECENTRALISED COMPUTERS
S-XCOM: SOFT. DEV.: FOR CROSS-SUPPORTED COMPUTERS
S-XMPU: SOFT. DEV.: FOR CROSS-SUPPORTED MICRO-COMPUTERS

T-T: TASK/TASK COMMUNICATION
TER: (REMOTE) TERMINAL
FAP: REMOTE FILE ACCESS PROTOCOL
DLL: DOWN LINE LOAD/ UP LINE DUMP
RJC: REMOTE JOB CONTROL
RPC: REMOTE PROCEDURE CALL
PAG: PAGING (VIRTUAL MEMORY)

 Y/N: SERVICE NEEDED
   N: NUMBER OF STATIONS
 ADR: AVERAGE DATA RATE PER STATION (BYTES/SECOND)
 PDR: PEAK DATA RATE (BYTES/SECOND)
 MAX: MAXIMUM DISTANCE
MTBF: DESIRED MTBF PER STATION (HOURS)
MTTR: DESIRED MTTR (HOURS)
 NAW: NIGHT OR WEEKEND SERVICE ESSENTIAL (Y/N)
```

TABLE 3 : SOFTWARE DEVELOPMENT

			N #STATIONS	ADR BYTES/S	PDR BYTES/S	MAX METRES	MTBF HOURS	MTTR HOURS	NAW SERVICE	COMMENTS
G-CENT	T-T	Y	100?		300K?	20K			Y	
	TER	Y	*	*	*	*	*	*	*	
	FAP	Y	*	*	*	*	*	*	*	
	DLL	N	-----	-----	-----	-----	-----	-----	-----	
	RJC	Y	*	*	*	*	*	*	*	
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	N	-----	-----	-----	-----	-----	-----	-----	
G-DECE	T-T	Y	100?		300K?	20K			Y	
	TER	Y	*	*	*	*	*	*	*	
	FAP	Y	*	*	*	*	*	*	*	
	DLL	N	-----	-----	-----	-----	-----	-----	-----	
	RJC	Y	*	*	*	*	*	*	*	
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	?			10M?					PAGING FOR WORKSTATIONS?
G-CERN	T-T	Y	100?		1M	20K			N	
	TER	Y	*	*	*	*	*	*	*	
	FAP	Y	*	*	*	*	*	*	*	
	DLL	Y	*	*	*	*	*	*	*	
	RJC	N	-----	-----	-----	-----	-----	-----	-----	
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	N	-----	-----	-----	-----	-----	-----	-----	
G-OUTS	T-T	N	-----	-----	-----	-----	-----	-----	N	for estimates, see E-OUTS
	TER	Y	50?							
	FAP	Y	*	*	*	*	*	*	*	
	DLL	Y	*	*	*	*	*	*	*	
	RJC	Y	*	*	*	*	*	*	*	
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	N	-----	-----	-----	-----	-----	-----	-----	

* MEANS 'INCLUDED IN OTHER ESTIMATES'
----- " 'NOT APPLICABLE'
? " 'DOUBTFUL - NEEDS FURTHER STUDY'
BLANK " 'UNKNOWN'

G-CENT: GENERAL COMPUTING: ACCESS TO CENTRALISED COMPUTERS
G-DECE: GENERAL COMPUTING: ACCESS TO DECENTRALISED COMPUTERS
G-CERN: GENERAL COMPUTING: TRAFFIC BETWEEN VARIOUS CENTRALISED AND DECENTRALISED COMPUTERS
G-OUTS: GENERAL COMPUTING: TRAFFIC TO OUTSIDE CERN

T-T: TASK/TASK COMMUNICATION
TER: (REMOTE) TERMINAL
FAP: REMOTE FILE ACCESS PROTOCOL
DLL: DOWN LINE LOAD/ UP LINE DUMP
RJC: REMOTE JOB CONTROL
RPC: REMOTE PROCEDURE CALL
PAG: PAGING (VIRTUAL MEMORY)

Y/N: SERVICE NEEDED
N: NUMBER OF STATIONS
ADR: AVERAGE DATA RATE PER STATION (BYTES/SECOND)
PDR: PEAK DATA RATE (BYTES/SECOND)
MAX: MAXIMUM DISTANCE
MTBF: DESIRED MTBF PER STATION (HOURS)
MTTR: DESIRED MTTR (HOURS)
NAW: NIGHT OR WEEKEND SERVICE ESSENTIAL (Y/N)

TABLE 4 : GENERAL COMPUTING

			N #STATIONS	ADR BYTES/S	PDR BYTES/S	MAX METRES	MTBF HOURS	MTTR HOURS	NAW SERVICE	COMMENTS
O-LOCL	T-T	N	-----	-----	-----	-----	-----	-----	N	ESTIMATES FOR 1 BLG. TOTAL BxBIGGER
	TER	Y	50		1M?	1K				
	FAP	Y								
	DLL	Y	*	*	*	*	*	*	*	
	RJC	Y	*	*	*	*	*	*	*	
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	?			10M?					PAGING FOR WORKSTATIONS?
O-CERN	T-T	N	-----	-----	-----	-----	-----	-----	N	FOR ESTIMATES, SEE G-CERN
	TER	Y								
	FAP	Y								
	DLL	N	-----	-----	-----	-----	-----	-----	-----	
	RJC	Y								
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	N	-----	-----	-----	-----	-----	-----	-----	
O-OUTS	T-T	N	-----	-----	-----	-----	-----	-----	N	FOR ESTIMATES, SEE E-OUTS
	TER	Y								
	FAP	Y								
	DLL	N	-----	-----	-----	-----	-----	-----	-----	
	RJC	Y								
	RPC	N	-----	-----	-----	-----	-----	-----	-----	
	PAG	N	-----	-----	-----	-----	-----	-----	-----	

* MEANS 'INCLUDED IN OTHER ESTIMATES'
----- " 'NOT APPLICABLE'
? " 'DOUBTFUL - NEEDS FURTHER STUDY'
BLANK " 'UNKNOWN'

O-LOCL: OFFICE AUTOMATION: LOCAL TRAFFIC (WITHIN DIVISION)
O-CERN: OFFICE AUTIMATION: INTER-DIVISION TRAFFICMB OR SMART)
O-OUTS: OFFICE AUTOMATION: TRAFFIC TO OUTSIDE CERN

T-T: TASK/TASK COMMUNICATION
TER: (REMOTE) TERMINAL
FAP: REMOTE FILE ACCESS PROTOCOL
DLL: DOWN LINE LOAD/ UP LINE DUMP
RJC: REMOTE JOB CONTROL
RPC: REMOTE PROCEDURE CALL
PAG: PAGING (VIRTUAL MEMORY)

Y/N: SERVICE NEEDED
N: NUMBER OF STATIONS
ADR: AVERAGE DATA RATE PER STATION (BYTES/SECOND)
PDR: PEAK DATA RATE (BYTES/SECOND)
MAX: MAXIMUM DISTANCE
MTBF: DESIRED MTBF PER STATION (HOURS)
MTTR: DESIRED MTTR (HOURS)
NAW: NIGHT OR WEEKEND SERVICE ESSENTIAL (Y/N)

TABLE 5 : OFFICE AUTOMATION

APPENDIX 7

Implementation Models for a Future Communications Programme at Cern

Author: WG2 SG3, C. Piney
Date: 3 Feb 83
Version: 11
Status: Final

The main recommendations of this subgroup are as follows: CERN should

1. provide a Backbone network which can be accessed from anywhere on the CERN site; develop the capabilities of CERNET to serve as this Backbone until such a time as it is superseded by other products;

2. select and support one local area network to be used as far as possible wherever one is required;

3. design, produce and maintain Gateways between the Backbone and a number of standard LANs;

4. develop and support the capability of using these facilities for transparent terminal access across the network;

5. provide support for devices obeying the X25 recommendations.

Provision of the necessary resources - capital outlay, manpower, maintenance - and their management must be accepted at a level commensurate with that of the main computer centre.

1. BASIC APPROACH

The period considered in this document covers the next ten years.

In order to simplify the problem posed by the large number of possible models, the set of all such models was first subdivided into categories. Each of these categories was then examined and evaluated. Finally, implementation models in the categories retained were then considered in detail.

Four main categories were identified as follows:

1. an extension of the current approach, i.e. CERNET plus INDEX (part 3);
2. heterogeneous local area networks interconnected on an ad hoc basis (part 4);
3. a global site-wide network providing no support for LANs (part 5);
4. a 'Backbone' network combining support for hosts and LANs (part 6).

The characteristics of each category with an assessment, will be presented in the following sections. The implications of the selection of a CERN standard LAN will also be examined. The assessment will be based on the essential requirements as laid down in T/W34. Only those requirements that appear to cause problems for a particular category will be considered in its assessment; those not mentioned can be assumed to be satisfied.

2. ESTIMATIONS OF THROUGHPUT

The figures below are included in order to give an idea of the magnitude of the problem.

It should be noted that the estimates given here agree fairly closely with those obtained from the survey carried out by SG2 (T/W33).

Each of the applications for which the estimations are made is critical to the future working of all computing at CERN. Firm decisions as to the techniques must be taken to allow design parameters to be fixed. If this is not done, the communication facilities may not be suited to the demands placed on them.

The estimations are based on the constraints implied by the data model described in T/W25. It should not be forgotten, however, that the effect of traffic patterns depends on the technology:

1. on a packet 'broadcast' system, all traffic on the network reduces the bandwidth available for all other users;
2. on a packet 'switched' system, traffic only affects the available bandwidth of users needing to traverse any of the switching hardware concerned in the original transfer.

2.1 FILE SERVER

The figures that follow include a reasonable margin to allow for future changes in the way in which people work. The figures are based on the assumption that many users will require access to a 'centralised' file base. This access will not be required for each manipulation, since a regional file-manipulation machine will perform that task. The central file server will used be as master back-up.

The main bulk of transfers will be made up of files. A typical file is of the order of 1 Mbits (i.e 4 000 lines). The average user will access or restore such a file on average every 2 hours. During peak hours, up to 5 000 simultaneous 'users' (terminals, computers, microprocessors, automatic monitors, etc.) must be catered for.

This estimate leads to an **average data rate of 1 Mbit/s.**

If, in addition, one assumes that a user should have 'instant' access to his file, the response should be of the order of 1/10 second. This implies a peak rate of 10 Mbps.

2.2 PRINT SERVER

Future printers will support not only character, but also digitised graphics, with a density of the order of 300 points per inch. Assuming that, on average, 10% of any page is taken up with graphical data (1 Mbit), compressed by a factor 10, plus 2000 printing characters, a page represents about 1/8 Mbit. A sustained rate of 1 Mbps would provide a printing rate, therefore, of 8 pages/second.

2.3 TERMINALS

There will increasingly be two classes of terminals:

1. those that are currently considered 'normal', i.e. character-oriented terminals running at up to 9.6 Kbps;

2. high-speed terminals, incorporating some degree of local intelligence and frequently transferring data in bursts of one line or one page or more.

There is, naturally, an overlap between these classes (see also T/W25) since, for example, line and even page-oriented terminals can frequently be satisfactorily supported over 9600 bps asynchronous lines.

For such terminals, the response should not exceed 1/10 second. Since a screen-full of characters is of the order of 2000 characters and the response may entail rewriting the whole screen, a rate of the order of **40 Kbps** is required to support the refresh of a full screen at an acceptable rate.

2.4 ELECTRONIC MAIL

This is yet another service that all connected users will require.

Although the choice of implementation of such a service has a considerable impact on the amount of mail data transferred, the total amount of such data does not put a significant extra load on the overall facilities.

2.5 VERY HIGH-SPEED REQUIREMENTS

There are certain very high-speed data-transfer requirements related to the operation of the central computers, such as specialised file access, interactive graphics applications, magnetic tape handling, etc. Such applications will need a special solution (e.g Hyperchannel as at LBL) and will not be considered as a factor in the CERN-wide throughput estimates.

In the same way, the high-speed transfer of large quantities of experimental data has not been considered, since such transfers will have to be handled by special methods, e.g. dedicated point-to-point links.

2.6 ADDITIONAL FACTORS

The other requirements, e.g. maximum allowed command-response delay for accelerator control, etc. must be added here to increase the precision of the figures. The requirement is currently assumed to be that, ignoring processing time in the responder, the delay imposed by the transmission system should not exceed 2/100 second for short communications (one packet of about 50 bytes as command, and the same as response).

Although further applications not detailed here will have some effect on the overall figures, it is not believed that the overall change will be much more than 20%, and the current document is working on a 'safety factor' which is considerably higher than this!

3. EXTENSION OF CURRENT APPROACH

3.1 EXTENSION INTO LEP

The extension into LEP would entail support over larger distances and provision for an increased number of users.

The ability to use the cables to be installed for LEP communications would also be desirable (see SG1 report T/W37).

3.2 ADDITION OF LANS

It has already been pointed out that LANs will be used both within experiments and for special applications (e.g. office automation). These will need to be connected. Studies are already being carried out into ways of providing connections into CERNET (T/W27, T/W30).

3.3 ASSESSMENT

For INDEX and CERNET, the main factors to be considered are distance, speed, reliability, number of users and price. These, plus other decisive criteria, will be examined below.

One point to be considered is that, for both CERNET and INDEX, the delays and marginal costs involved in connecting new users are high compared to what is foreseen for projected future systems.

3.3.1 INDEX

1. distance, as such, is no problem;

2. there is no software development, and therefore no need for programming staff on-site; there is also no need for support from the central computer operators;

3. the size of CERN, the need to avoid installing very long cables and the requirement, in some cases, for 'localised' communications lead to the use of PACXs other than those in the computer centre, and thereby increases, for some users, the number of PACXs traversed, which, in turn, increases the problems due to jitter. This 'jitter' problem will be solved by the Rev. E microprocessor from Gandalf which provides the capability of synchronising all PACXs to one master clock;

4. the 'multi-star' approach of Index can be contrasted to a multi-drop approach:

   a) the number of cable-runs and length of cable is greater for Index then for a multi-drop system, whereas the cable and connectors may well be cheaper for Index;

   b) the multi-drop approach allows easy installation of new connections, but only within regions already catered for in the original planning and installation of the basic cable;

5. the speeds offered are sufficient for the majority of terminals currently installed at CERN but cannot satisfy the requirements of high-speed devices;

6. one inherent characteristic of circuit-switched systems is that no checking is done on the integrity of the transferred data. It is up to the end-users to perform such checks if necessary, given the very low error-rate of the system;

7. the addressing scheme has 2 shortcomings:

   a) the range is not sufficient to cover the expected number of hosts;

   b) the destination address depends on the PACX to which the terminal is connected, since the route has to be given.

   The first problem can be solved by the use of alphanumeric destination names. This solution would allow the inclusion of about 100 hosts per PACX (i.e. about 500 at present). For the second problem, by the Summer, Gandalf expect to release a new version of the PACX code which will support through-PACX routing;

8. the asymmetry of INDEX, where ports are defined as either 'host' or 'terminal' ports, is inappropriate for sophisticated devices, such as workstations, which need to act as both hosts and terminals;

9. the way Cern uses INDEX does not allow terminals to remain addressable when not connected to a host, and does not allow 'talk mode' between terminals. These capabilities could be provided by connecting terminals in a different way to the PACX (e.g. direct V24 connections from terminals to the AMTB card in the PACX).

3.3.2 CERNET

1. To communicate over the distances required, Cernet can either adopt the approach outlined in T/W37, or use repeaters (every 1500 m), or slow down the transmission speed;

2. some Cernet hardware and software is produced and maintained at Cern; in addition, the operation of the system entails some effort from the central computer operators;

3. the speed provided by CERNET will not be adequate for the total demand envisaged for the file server, but is sufficient for all current applications;

4. additional users could be catered for by the addition of extra switching nodes, and the developments mentioned above (e.g. T/W30) could make this simpler;

5. all data transferred is checked for integrity, and end-to-end error recovery procedures are available;

6. current CERNET hardware connection costs are about 10 000 SFR, which is higher than many of the objects likely to need connection - it is likely that this can be reduced significantly by further developments (T/W27, T/W30);

7. workstations, microprocessors and 'personal computers' will have to be supported, either directly or indirectly, thereby adding, over the next ten years, up to 5000 more addressable entities thus exceeding the CERNET addressing capabilities;

8. the current mean load on CERNET - at least on the IBM front-end - is of the order of 25%, and the factor 5 improvement required to reduce this to the order of 5% (as specified in 'Essential Requirements: T/W34) does not appear likely without a different approach (e.g several paths into the 'main' machines);

9. CERNET obeys no international standards at any level.

3.4 CONCLUSIONS

1. Index will have increasing difficulty in satisfying the demands made on it by the terminal population.

2. The traffic patterns envisaged will ultimately put an excessive load on CERNET.

3. Installation costs and delays for both systems are high relative to what can be expected from future systems.

4. The requirement of most subscribers to have both INDEX and CERNET connections increases the cost in terms of money and manpower.

The points given above indicate that both INDEX and CERNET will run out of capacity during the period envisaged, but can both provide most of the services required during the initial period.

4. AD HOC LAN INTERCONNECTION

This category has been retained for consideration mainly because the arguments for its rejection help to underline some of the main aims to be borne in mind. The disadvantages of this approach are:

1. no guarantee, for any user, of concerted support;

2. connectivity, i.e. ability of one user to communicate with another, not guaranteed, since this could entail routing traffic through a number of LANs before it arrives at its destination. In addition, the number of Gateways required under such conditions increases as the square of the number of different types of LAN;

3. capacity of the system uncertain and possibly dependent on the path taken;

4. impossibility of organising a rational site-wide cabling strategy;

5. access to the central services uncertain.

In addition to these reservations, it should be noted that for there to be any common facilities, there must be a common and preferably shared set of protocols for their use. This is much more difficult for this category of model.

5. ONE CERN-WIDE NETWORK

This category is the converse of the previous one. It will ensure homogeneity by drastically restricting variety. It could never be adopted in any organisation such as CERN which aims to provide support to a varied population of users, and intends not to drop too far behind modern developments.

Moreover, much of the traffic at CERN is localised to groups, buildings, etc. and this way of working will increase with the advent of LANs, file-servers, etc. Under these conditions, it is much more efficient not to load a CERN-wide network with traffic that does not require CERN-wide services.

6. THE BACKBONE APPROACH

6.1 THE MODEL

The Backbone category of models aims to simplify the problem of connection by insisting that LANs intercommunicate only across the Backbone. The Gatewaying problem - at least up to the protocol level supported across the Backbone -, for any LAN to be connected, is then reduced to connecting this LAN to the Backbone.

General services will then be available to any subscriber since both the service and the subscriber are connected either directly to the Backbone, or indirectly via a LAN.

To be generally useful, the Backbone must therefore reach into, or be reachable from, any part of the site. This access to the Backbone should, obviously, also satisfy the essential requirements with respect to price, throughput, response, reliability, etc.

6.2 ASSESSMENT

6.2.1 Essential Requirements

Given the characteristics outlined above, the following assessment of this category can be made: to satisfy the number and price requirements for the connection of terminals, it is likely that terminals will need to be connected indirectly - via either concentrator boxes or LANs. With this proviso, none of the other essential requirements is ruled out in this category of model.

6.2.2 Terminals over Packet Networks

Since the backbone model implies carrying terminal traffic over a packet network, the following disadvantages of such an approach should be borne in mind:

1. unevenness of screen output due to stochastic arrival of packets;

2. some combinations of terminals and applications lead to the use of single-character packets to ensure full transparency, which entails considerable network overheads;

3. treatment of 'special' keys (e.g. break) can be fairly complex.

It is, however, believed that these drawbacks can be minimised in a well-designed system, since, even in some currently available systems, they are virtually unnoticeable.

7. IMPLEMENTATION DECISIONS

7.1 SELECTION

The assessments of the four categories of model leads to the rejection of the two which either violate the essential requirements, or are too restrictive for the CERN environment:

1. the ad hoc interconnection of LANs;

2. the unique CERN-wide network.

The two categories retained are not totally different from each other and can be considered as two points on a smooth development path:

1. the extension of the current approach;

2. the Backbone approach.

7.2 TRANSITION SCENARIO

When a suitable Backbone is available, it should be integrated with CERNET and INDEX as smoothly as possible. This entails planning future developments of these services with the intention of simplifying the future changeover.

Today, CERNET offers a better Backbone service for CERN than could be provided by any commercial system currently available; for some time to come, CERNET will also remain competitive with local area networks. INDEX provides a service that will remain valuable in certain areas for many years to come.

The installation of any future system should be planned so as to preserve, as far as possible, the investment so far made in both CERNET and INDEX.

7.2.1 LANs and the CLAN

In order to limit the proliferation of different LANs:

1. CERN will select one local area network as its preferred and recommended LAN ("the CLAN");
2. full hardware and software support will be provided for the CLAN;
3. CLANs will gradually be installed throughout CERN as the standard means of connection to the general CERN communication facilities;
4. a number of other LANs will be selected for a reduced level of hardware and software support. Examples of such special LANs (SLANs) could be the digital PABX, the Hyperchannel, etc..

7.2.2 CERNET and the Backbone

The following steps form a possible scenario:

1. CERNET will act as the Backbone: in addition to linking many host machines, as it does today, it will also serve to interconnect the supported LANs;
2. subscriber machines will be moved, where possible, from CERNET onto LANs;
3. the Backbone network will be selected and installed;
4. CERNET will be Gatewayed into the Backbone at more than one point, so that immediate relief can be obtained for lines which are overloaded;
5. CERN supported LANs will then be removed from CERNET and Gatewayed directly onto the Backbone;
6. once the Backbone is reliable, redundant CERNET links will be removed. However, local area CERNETs will continue as long as they are competitive with new systems;
7. all new connections will then be made either via LANs or directly to the Backbone, and CERNET development will be frozen.

7.2.3 INDEX and LANs

In order to optimise the use of cables and computer ports, as well as integrating the various communication services, the following steps could be followed:

1. a Gateway will be produced to enable computers which are on LANs to access INDEX;

2. facilities will also be developed to allow all connected devices to be accessed across the network by terminals. There are two ways in which this will be done:

   - initially, by means of a PAD external to the subscriber, interfacing to the subscriber like a terminal and to the network like a subscriber;

   - ultimately, by extending the software in the subscriber to implement an agreed network terminal protocol;

3. whenever convenient, future connections will then be made to the nearest available CLAN.

7.3 GATEWAYING

Of paramount importance is the ability to interwork between the various objects connected to the projected communication facility. The ease and even feasibility of such interworking depends entirely on the effectiveness of the Gateways provided.

Only Gateways at the 'lowest' level (i.e. datagram level) are considered here.

The Gateway requirements may impact on the networks to be connected and the Gateway studies should therefore be started as soon as possible.

It is strongly suggested that, at least initially, the overall Gateway hardware be designed and produced at CERN, (making as much use as possible of commercially available equipment), since the design and implementation of the Gateways will affect the entire communication facility.

It is imperative that such Gateway developments be carried out with the aim of encouraging the use of international standards wherever possible. In this way, the networks installed at CERN will be able to adopt the various relevant international standards as they become available, with a minimum of disruption to existing users.

8. RECOMMENDATIONS

1. Development towards a Backbone capable of supporting the estimated throughput: i.e. a mean rate of 1 Mbps and a peak rate of 10 Mbps. This means an installed capacity of at least 10 Mbps, extendible by a factor of the order of 10 without a total change of technology.

   - Backbone accessible from any point on the site;

   - Smooth degradation of the Backbone service under heavy load.

   This Backbone could be implemented either by an improvement of existing CERNET technology, by a commercially available product (for which there seems no contender at present), or by a new CERN development. Some additional technical considerations are given in appendix A. It is recommended that **by 1984, the final choice as to the Backbone technology must have been made.**

2. Selection of a 'CERN standard LAN' (CLAN) which will be the normal means by which equipment is connected. All CERN buildings should eventually be wired for this CLAN, and it should serve as the standard, though not the only, means of access to the CERN data-communication facilities. The selection of this CLAN will have to be carried out in close collaboration with all the people affected (physicists, everyone concerned with accelerators, programmers, administrative staff, etc.). Possibly reduced support for other 'standard' LANs. For these, the connection of devices, for example, would be the responsibility of the user.

3. Installation of Gateways on the Backbone, to the CLAN and to other 'standard' LANs.

   - Gateways offering a throughput comparable to that of the networks they interconnect, i.e. not being a bottleneck in the communication (although obviously adding a delay).

   - Packet transit time (LAN - G/W - Backbone - G/W - LAN) of the order of better then 1/10 second.

4. Development of capabilities for transparent terminal access across these facilities. A first step towards this aim would be the development of a Gateway from either a CLAN, or the Backbone, to INDEX.

5. Support for devices obeying the X25 recommendations. This entails, effectively, setting up an X25 'envelope' to the communications facilities:

   - access to DTEs anywhere within or beyond CERN on an X25 network;

   - X25 access by DTEs on the CERN site;

   - X25 access provided by a special connection into the nearest CLAN.

9. COST ESTIMATES

This section presents the typical costs of one possible hardware implementation of the proposed solution. This is intended to indicate what sort of costs can be expected if a satisfactory system is to be provided. The costs have mostly been derived by referring to existing, commercially available products. In some cases, however, few products exist. In other cases, the available products do not map very well on to the CERN requirements. Furthermore, the influence of evolving technology on prices is uncertain. These costs should therefore be treated only as a very rough guide.

1. Backbone:
   This is particularly difficult to cost, because there is no commercially available product which would meet the essential requirements. However, taking the cost of CERNET and ETHERNET prices as guidelines, a rough estimate of the cost of each connection would be 5000 SF, for the basic hardware, plus about the same again for installation and a share of the cost of the infrastructure. Thus, for a Backbone with about 200 connections, the cost would be about 2 MSF. The manpower required for development and installation of the backbone will depend on the choice of system and is a subject for future study. The staff requirement, after initial development, would be about 5 people full time. This is similar to the staff now required for CERNET maintenance, expansion, and development.

2. CLAN:
   For the selected CLAN to meet the essential requirements, the connection cost must be kept low. Initial estimates, based on currently available products, suggest a connection cost per subscriber computer in the range 1000 to 2000 SF. (N.B. Terminal connections are considered separately below). The share per device of the infrastructure would be about 500 to 1000 SF, so the total cost per connected device would be in the range 1500 to 3000 SF. The total cost for 1000 connections would thus be between 1 and 3 MSF. The manpower required for maintenance, expansion, and development would again be in the region of 5 people.

3. Terminal interfacing:
   There are two approaches, both of which will be needed:

   - hardware terminal interface:
     taking the single-user triple-X PAD developed at ISR for UTI-NET as an example, the cost per terminal should not be more than 500 SF. So, for example, the connection, in this way, of 2000 terminals, would cost 1 MSF;

   - subscriber software package:
     this entails the inclusion in the host operating system of a 'network terminal protocol'. It should be noted that modern operating systems are designed to allow such an approach, but there are still many systems on which this is extremely difficult. Estimates for the provision of a basic service must therefore range from 6 man-weeks in the best case to 6 man-months in the worst case.

4. Gateways:

The hardware price is only for the 'Gateway' part, and does not include the interfaces to the two networks to be connected, as these costs are included in the connection costs for these networks. In this way, the cost is, to a large extent, independent of the networks that are Gatewayed.

Initial estimates are of 500 SF for the hardware.

The expected manpower requirements for the production of a satisfactory Gateway, once the required hardware is available is

- of the order of 4 man-years for the first Gateway produced. This Gateway would be so designed as to serve as the building-block for all future such Gateways;
- 9 man-months to provide a new Gateway;
- one programmer half-time for the run-in period for each Gateway (about 1 year).

5. X25 equipment:
The hardware cost, based on currently available equipment and present trends, is estimated at 1000 SF per connection; the software effort would be of the order of 1 man-year to set up such a system.

6. Specialized high performance LAN for the computer centre:
There will probably be a need for a very high performance LAN in the computer centre, interconnecting large mainframes, file servers, file backup systems, and so on. Taking Hyperchannel from NSC as an example, and assuming ultimately 10 devices are connected, the cost would be about 1 MSF. The manpower requirement would probably be about two people, but it would probably be best to combine the support for the special high performance LAN and the Backbone.

10. CONCLUSIONS

The data communication needs of CERN over the next decade or more are best provided by a two-level approach. A "Backbone" network will be required at the higher level; this will interconnect many different types of devices but will be used in particular to interconnect local area networks.

A particular local area network implementation (the CLAN) should be agreed by CERN and its users; many of these CLANs will have to be installed around the site, and they will provide the primary data communications interface.

A limited number of types of specialised local area networks (SLANs) should be chosen and supported by CERN. The Gateway connecting a SLAN to the Backbone network should be supported by CERN, but with an overall level of support that may not be as complete as for the CLAN.

A specific SLAN will have to be installed in the computer center to interconnect the large mainframes, certain file storage devices, and other peripherals; this SLAN also should be connected to the Backbone network by a CERN-supported Gateway.

Appendix A

BACKBONE NETWORK IMPLEMENTATIONS

This section presents several ways in which the services of the Backbone network can be provided.

A.1 HOMOGENEOUS SITE-WIDE NETWORK

This approach entails the selection and setting up of one network satisfying the requirements outlined earlier and spanning the entire site. This puts very serious restrictions on the technology that can be selected, since, for example, CSMA/CD medium-access techniques cannot be used for high-speed communications over such distances.

The only remaining technologies seem to be the ring, the token bus and the classic mesh-structure. Each of these warrant further study before a decision can be taken as to their final applicability.

A.2 CENTRALISED SWITCH

The requirement that the Backbone be reachable from anywhere on the CERN site does not imply that the Backbone should be present everywhere on the site. The objects needing to connect to the Backbone could access it, if necessary, down long point-to-point links. In this case, the technology used for the Backbone would not be constrained by the large CERN distances, and CSMA type techniques would once again be valid.

The attraction of this approach is that it eliminates the distance restriction which seems to cause considerable problems for very high speeds with most technologies. The drawback would be the ad hoc approach that would need to be adopted, over large distances, for connecting the objects to the Backbone.

A.3 SEPARATION OF TRAFFIC

This approach would help to overcome the drawbacks of the two previous suggestions:

1. installation of a rapid, short-distance network, to which all high-speed objects and servers will be connected directly (high price);

2. provision of an access network to this 'nerve-centre' by means of a technology offering perhaps lower speeds but capable of covering greater distances. In this case, the connections of this access network to the nerve-centre would still be costly, but the connections to the access network itself should be considerably cheaper.

This model reduces, in fact, to the case of the Backbone with special LANs connected, and has therefore already been treated. It serves, however, to show how such special LANs have a role to play in reducing the overall load on the Backbone.

Appendix B

STEPS TOWARDS A SMOOTH TRANSITION

To ensure a smooth transition from the current situation, it is proposed that the following steps be taken:

1. study methods of improving throughput possibilities of CERNET;

2. develop medium-speed concentrators (4.8-64 Kbps) for connecting terminals and similar devices (e.g. microprocessors) to CERNET;

3. study ways of providing for the physical connection of > 400 objects;

4. define a CERN-wide 'Gateway-level' addressing scheme which should remain valid for in the foreseeable future, i.e. which will allow direct addressing of at least 8000 subscribers, and will allow addressing of 400 or so LANs each with several hundred subscribers;

5. extend the addressing scheme to allow consistent communication with subscribers outside CERN;

6. convert existing CERNET protocols to support this addressing scheme.

CONTENTS

APPENDIX D
TO THE LEP COMPUTING PLANNING GROUP REPORT

WG3/cmj/82/53 Rev. 5
17th March 1983

Final Report of Working Group 3 on General Computing Services and Facilities for the LEP Era

(The members of Working Group 3 are given in appendix A)

CONTENTS

1. MAJOR CONCLUSIONS AND RECOMMENDATIONS

The major conclusions and recommendations of WG3 are summarised below. Further recommendations on a number of detailed matters are included in the text of the report.

1. The next decade will require not only the strengthening of the central batch processing power, but also a major improvement of the human interface to all types of computing service, that will enable people to achieve more, and to work more efficiently. (See ch.3, sections 5.2, 5.3, and 5.4).

2. Based on estimates of the likely computing load from LEP and other activities, WG3 **recommends** that the central batch processing power be expanded by at least a factor 2 by 1987, in time for the start of LEP physics. This assumes that the 1/3:2/3 inside:outside CERN policy will be enforced. At our open meetings, and in private communications, many physicists have emphasised that this factor 2 is a minimum figure, and no-one has suggested less. Extrapolations from the actual usage over the past decade suggest a factor of 4, (see chapter 3, section 4.1 and figure 5).

3. This power may be installed in an incremental manner when needed, as explained in the report. However, it is vital that the power be installed in good time to avoid periods of saturation. Failure to do so would lead to a tremendous waste of a precious resource in the next era, namely human effort. (See sections 5.3, and 9.1).

4. The major change in style in the next decade must be an improvement of the human interface, an area largely neglected by manufacturers in the infancy of computing. This improvement will come mainly from **interactive computing,** a subject which is just coming of age. (See section 5.4, and chapter 6).

5. Taking a long term view, we consider the most effective way to achieve this will be through Personal Workstation Computers (PWSs) and local area networks (LANs), connected via a broad spine to the computer centre. Whilst it is too early to purchase PWSs in large numbers today,

WG3 **recommends** the purchase of small numbers immediately, and the creation of a PWS Project. This project should put together a serious programme to investigate how such stations can best be used and interfaced in an overall integrated network. This practical knowledge cannot be purchased from the manufacturers. (See chapter 6).

6. The increasing success of super-mini computers such as VAX, for general usage within those physics groups who are able to purchase them, cannot be ignored. WG3 **recommends** that CERN provides support for them in the form of libraries, and packages to interface them to the network, and to the computer centre. CERN must plan for their needs of file-base support and communications. However WG3 **does not recommend** the installation of multiple (of the order 10/15) VAX computers, running the VMS operating system, as a way of providing the **general** interactive computing service in the computer centre, despite their success in the above **private** environments. (See chapter 6).

7. We **recommend** that the general interactive service needs be filled in the short term on the central computers. In particular, subject to a successful technical evaluation and trial, we **recommend** the installation of the interactive system VM/CMS on the central IBM compatible computers, initially as a replacement of TSO. This is the most cost-effective way to satisfy the general interactive need now, and the needs of the occasional user later. Nevertheless, users having heavy interactive needs should start soon with the PWSs. (See chapter 6).

8. The enormous investment that has already been made in WYLBUR, especially in EXEC files, requires that the WYLBUR service be maintained roughly at the present service level for many years to come, at least until a VM/CMS service is well established, and has superior offerings. We **recommend** the minimum possible further development of WYLBUR features, but the WYLBUR service must not suffer degradation through the installation of VM/CMS. This will imply gradual addition of new hardware. (See section 9.1).

9. There are some very real dangers in this era of distributed computing. A proliferation of small and different systems could bring, (and is already starting to bring), an intolerable load on staff for support,

interfacing, and integration at the level of end-user services. It is essential to find mechanisms to contain the proliferation, and to ensure **integrated distributed computing.** (See chapter 5).

10. WG3 **recommends** central control of acquisitions of all such systems. We **recommend** early choices of standards for minis and PWSs in order to persuade groups who are not under CERN's control to follow them. (See section 5.4).

11. In general, major peripherals should continue to be sited in the computer centre for reasons of economy, sharing and early availability of new devices. This requires excellent networking, and integrated services, if they are to be easily and generally available to the distributed computers.

12. The IBM mass store is the most important of these peripherals. It requires expansion in volume and bandwidth on a short time-scale. Although one would ideally replace it with a more modern technology if it became available, it is unlikely that we can wait long enough to avoid upgrading the present device. We **recommend** the upgrade to an A3 model in 1983. (See sections 5.3.4 and 9.1.2).

13. The present role of the mass store as a central file server will increase in importance. We **recommend** that a detailed technical study of the implications be made. This should consider the justification for a specific file server computer for the network, as well as a network-wide filing system and the file backup requirements of distributed disks on site.

14. We **recommend** that a further technical study be made of a high speed connection between the many computers in the centre. This must decide between the future CERN backbone network and commercial offerings.

15. The present on-site training given in computing is far too limited. It is clear that, today, we are fulfilling the needs of neither the non-physicist users, nor the physicists, nor the computer professionals who should keep up with the rapidly changing outside world of computing. The conclusions both of the Software Workshop, (refs. 2 and 3), and the

CERN Joint Training Board support these remarks. WG3 **recommends** both an increase in the training and support given to all types of users, and the creation of a regular professional training programme. (See chap. 7).

2. THE MANDATE OF WG3 AND THE STRUCTURE OF THIS REPORT

Working Group 3 was established in May 1982 by the LEP Computing Planning Group (Chairman E.Gabathuler, CERN Directorate member responsible for computing). Its mandate was "to study and to make recommendations as to how to ensure adequate general computing services and facilities for LEP experiments, and consequently for the rest of the CERN laboratory". They were required "to report to the Planning Group by late Autumn", and "the time frame of the study should extend to beyond 1990". "The studies should bring out the costs, manpower and alternatives in order to show the scale of the problem". The working group has held 20 meetings, plus 2 full-day open meetings where a number of other people were able to contribute their ideas to our thinking.

All the documents of WG3 are filed under the CERN WYLBUR account CM.WG3. Versions suitable for direct printing are filed, with minor exceptions, in the library (PDS) file CM.WG3.PRINTER. The SHOw DIRectory command will give a list of available documents, which may then be printed via the LISt OFFline command. (See appendix C).

The structure of this report is summarised on the cover page, and a detailed table of contents is given at the end. Some words of explanation are appropriate. The recommendations are given at the beginning for obvious reasons. The next chapter covers the implications for LEP computing coming from the experiments themselves, and contains both an estimate of the processing capacity required, as well as a survey of other problems to be encountered. The next chapter on the present general computing services and facilities has been added as a consequence of suggestions made about the WG3 Interim Report. Although it delays the meat of the discussion on the future, it is felt necessary to illustrate the scale of today's problems.

The options and some likely models for the LEP era are discussed in the next chapter, and continued in the following chapter, which tackles the specific problem of how to provide for interactive computing. Training is treated in a chapter of its own because of its importance. A number of shorter topics have been grouped together in the next chapter, and lastly we consider the costs and manpower implications of our choices.

The names of those who have contributed to the Working Group are given in appendix A. Since the world of computing has seen fit to shower us with abbreviations unintelligible to the non-specialist, we have included a glossary as appendix B. We have completed the report with a full index of the working papers of WG3 (appendix C), and a list of references to other works.

3. IMPLICATIONS FOR LEP COMPUTING FROM LEP EXPERIMENTS

The purpose of this chapter is to set the **scale** of the likely data taking and analysis loads from LEP experiments, to point out some of the aspects of computing in a LEP (or any large) collaboration, and finally to indicate how these factors impact on the European HEP computing environment for the late 80's.

Accordingly this chapter is divided into four sections:

1. Physics assumptions

2. Data taking and analysis loads

3. Organisation of computing in LEP experiments

4. Consequences for European HEP computing.

3.1. Physics assumptions

a) The Z^o exists at a mass (≈100 Gev) accessible by LEP phase 1. The cross-section at the Z^o peak is 40 10^{-33} cm^{-2}.

b) The **time-averaged** luminosity, <L>, at the Z^o peak will be 1. 10^{31} cm^{-2} sec^{-1}.
Note that this is a factor 5 **below** the "pink" book design figure.

c) The major part of running in the first 2 to 3 years will be at the Z^o peak.

d) Data taking will be scheduled for 4000 hours per year.

e) The events will be high multiplicity (i.e. "complex").

The above assumptions will be used for estimates of data rates. It should be noted that all backgrounds are ignored, which could be a drastic assumption. Furthermore it is assumed that the two-photon process, which has a cross-section ≈10 times greater than the Z^o, can be rejected efficiently on-line if desired, and so will not contribute to the data written to tape.

3.2. Data taking and analysis loads

The above physics assumptions give an event rate per experiment of 0.4 Z^o events per second, or ≈ 1.5 10^6 Z^o events per 1000 hours.

3.2.1. Event size

The LEP experiments currently being considered have > 200 000 channels of electronics. It is difficult at this stage to know what the average event size will be. It is also quite possible that the use of powerful micro-processors on-line could eventually reduce significantly the event size. For the purposes of argument we shall assume the average event size to be 100K bytes (1 byte = 8 bits).

3.2.2. Recording medium

At 100K bytes event size, the average recording rate required for Z^o events is 40K bytes per second. Present day 6250 bpi 125 ips tape technology implies a **maximum** transfer rate of ≈750K bytes per second. The total capacity of a 2400 ft 6250 bpi tape is ≈150M bytes, allowing ≈15% for inter-block gaps. So using 6250 bpi tapes we have: A luminosity of $1.\ 10^{31}$ with an event **rate** of 0.4 per sec., and an event **size** of 100K bytes,
==> 1 6250 bpi tape written per hour.

One tape per hour is a very comfortable rate. However a factor of three increase in each of luminosity and event size would imply ≈10 tapes per hour, which is barely practical. Similarly, inefficient rejection of background could upset these numbers.

Of course recording technology does not stand still, and there are at least the following possibilities:

- Higher performance tapes and tape units: Progress in tape technology has been very rapid. The evolution from the 200 bpi 7-track tape to the modern 6250 bpi 9-track tape took less than 20 years. Considerably higher densities are technically **possible,** but no advance on the 6250 bpi tape is widely commercially available, although it is expected that IBM will announce a factor of ≈3 upgrade in the next year or so. However it took about 10 years for the 6250 bpi tape to "reach the counting-room". Hence it would seem rash to assume any tape technology beyond the present 6250 bpi tape for the LEP counting-rooms.
- Cheap robust large disks: One can imagine something like a 5 gigabyte (≈ 30 tapes) disk to act as a buffer store for a day's events. This would be convenient for fast quality checking and selective analysis of the data. For permanent storage the data would have to be copied off to tape. The gigabyte disk is only now available on large main-frame machines, so again this does not seem realistic for LEP experiments at the end of the 80's.
- Video-tape: This gives a huge capacity factor over normal tape, (recording rate similar to normal tape), but the hardware units for digital data are not commercially available. Hence this is impractical as the permanent storage medium, whilst disk looks to be the better buffer store device for the future.

- Video-disk: Single write, multi-read video disks of capacity of the order of hundreds of Mbytes per surface, and transfer rate of the order of 1 Mbyte per second, (using solid-state rather than gas lasers), can be expected in the next couple of years. In the long term this is a possible replacement for tape as the permanent storage medium. For the moment we can only await developments.
- Recording data in the Computer Centre rather than the counting-room: This is a more radical proposal for CERN, though it was the system used at Daresbury, and is presently in use at DESY. It is too large a subject for any detailed discussion here. Suffice to remark that it requires high quality links between the counting-room and the computer centre, plus the allocation of some fraction of centre resources to data recording. Our feeling is that there has been no **strong** demand for this from CERN experiments so far, and so we will not consider it further here.

So we shall assume that LEP experiments will record data **in the counting-room** on **present technology 6250 bpi tape.**

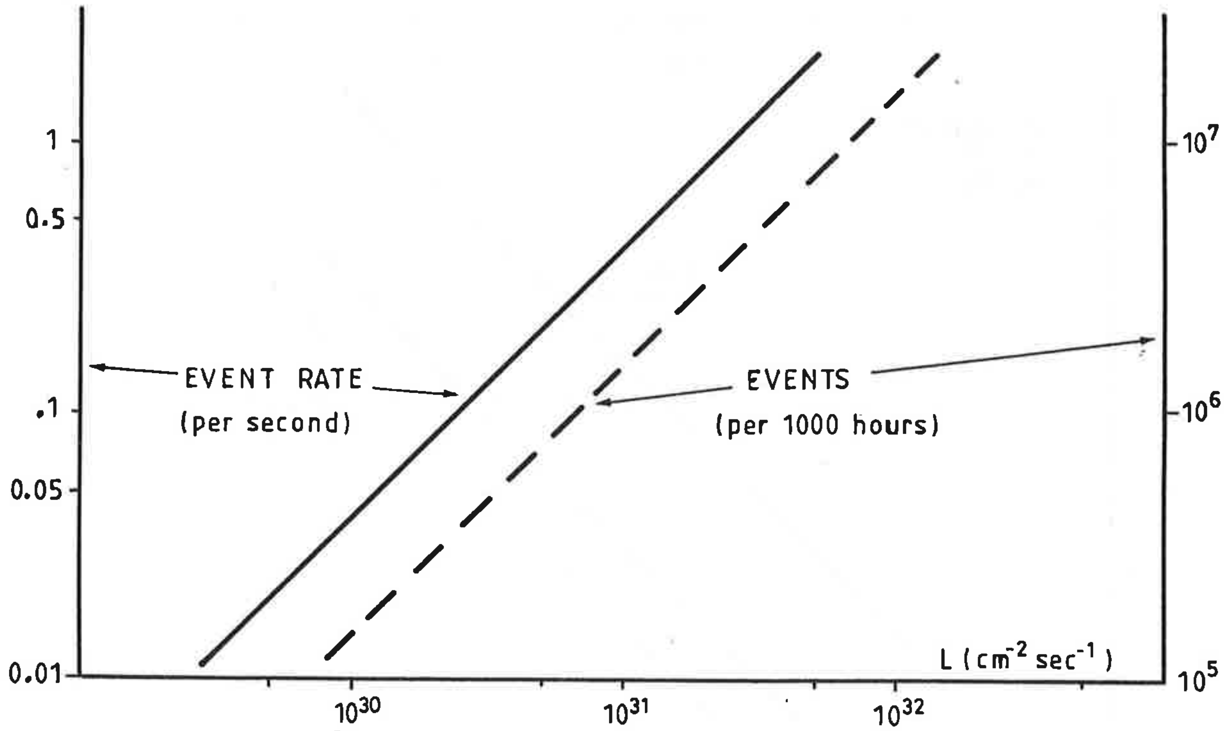

Figure 1: Z^0 event rate as a function of Luminosity, L.

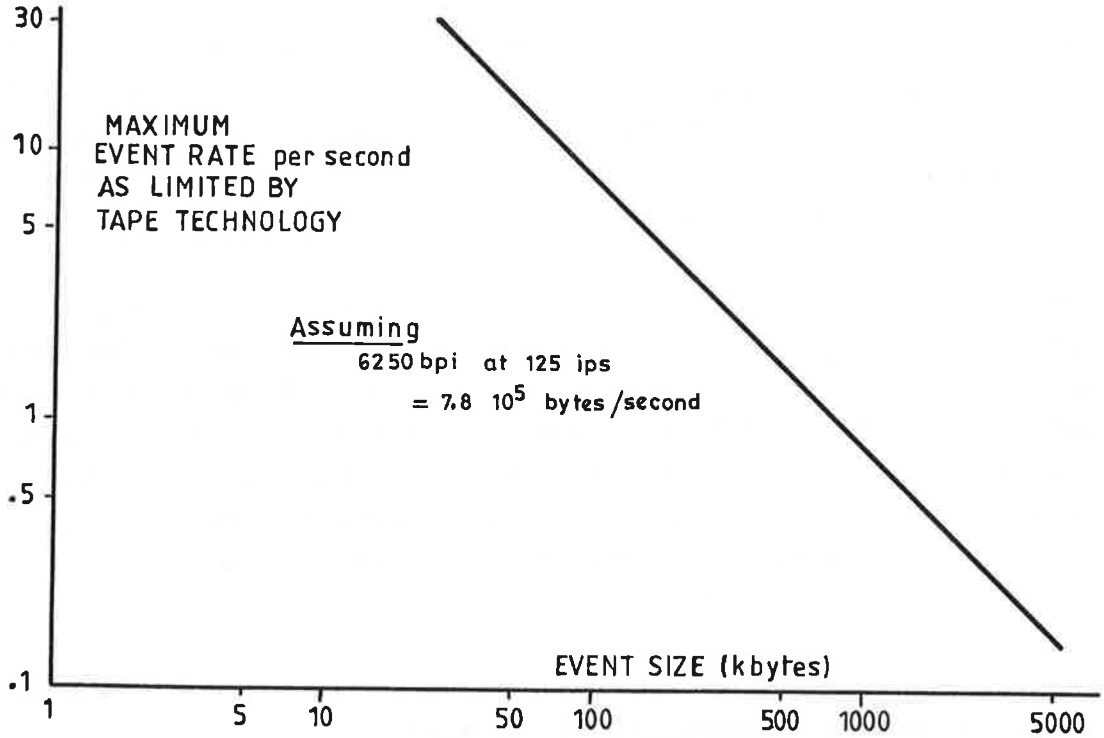

Figure 2: Maximum event rate, as limited by tape technology only, as a function of event size.

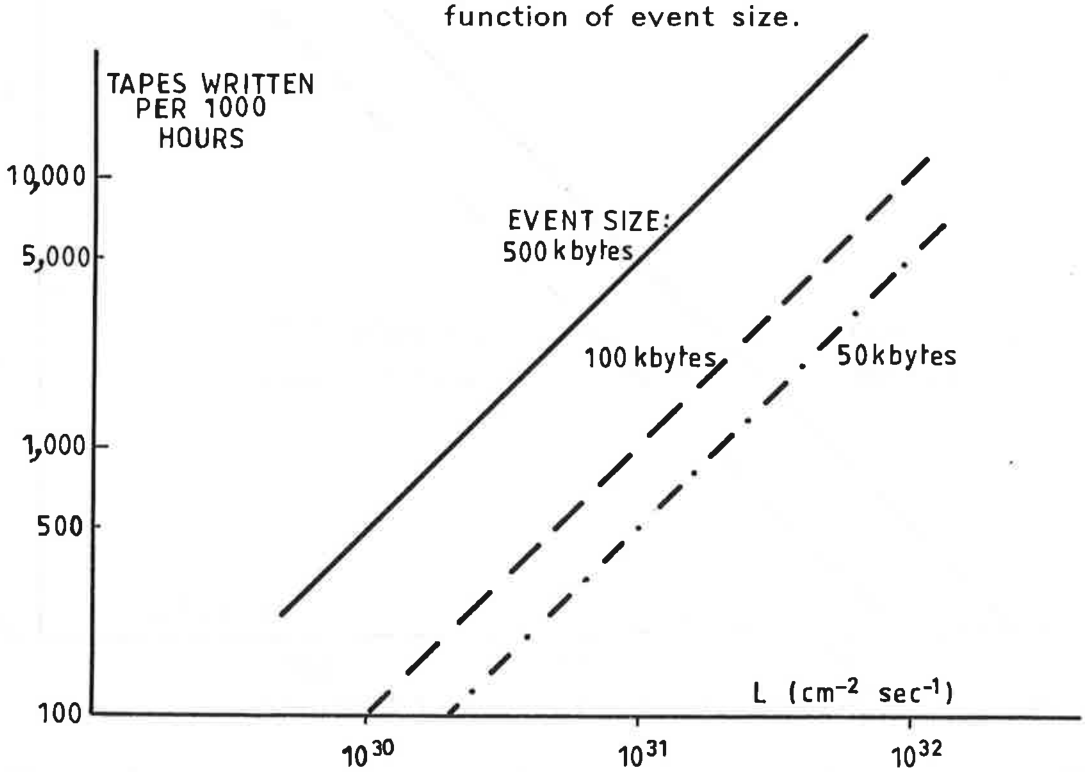

Figure 3: Tapes written per 1000 hours as a function of Luminosity for the different event sizes.

The interplay between L, event size, and number of tapes written is summarised graphically in figs. 1, 2, and 3.

3.2.3. Off-line analysis load

4000 hours of data taking per year leads, with the above assumptions, to each LEP experiment writing 4000 tapes containing ≈6. 10^6 events per year.

The off-line cpu analysis time needed is of course difficult to estimate with any confidence. We shall assume that **full** reconstruction (i.e. from raw data tape to 4-vectors) will require 20 seconds IBM 168 per event. Hence a year's data taking implies a cpu load per experiment of:

$6\ 10^6 * 20$ sec. (IBM 168)

≈ 30 000 hours ≈ 4-5 IBM 168's

Of course there are several factors which can change this estimate **significantly.** It is **reduced** by:

- LEP machine inefficiency:
- Experiment down-time, set-up time etc.:
- Not fully reconstructing all events:

It is **increased** by:

- Higher luminosity:
- Other computing needs: e.g. Calibration runs, Monte-Carlo, DST analysis, program preparation etc.

Several of the factors just mentioned are potentially large, so even the uncertainty on the estimate of 4-5 IBM 168's is hard to estimate! An uncertainty factor of 2 seems plausible, and experience suggests that the estimate will turn out to be too small rather than too large.

Finally the various assumptions leading to the estimates made in this section are summarised in Table 1

Table 1: Summary of Assumptions

Z^o	Mass	≈100 Gev
	cross-section	40 nanobarns
Time-averaged luminosity		$1 \cdot 10^{31}$ cm^{-2} sec^{-1}
=======>	event rate	0.4 sec^{-1}
		1.5 10^6 per 1000 hours
average event size		100K bytes
tape capacity		150M bytes
=======>		1 tape written per hour
Data taking scheduled for		4000 hours per year
Live time for 'good' data		4000 hours per year
======> good data per year		4000 tapes≈ 6 10^6 events
Fraction of events fully analysed		1
cpu time for full analysis		20 sec IBM 168 per event
======> cpu time per year		30 000 hours IBM 168
Uncertainty factor		≈2

3.3. Organisation of computing in LEP experiments

LEP experiments will be collaborations of upwards of 200 physicists from several research centres and universities, mainly but not exclusively from Western Europe. At any time a significant fraction of the participating physicists will **not** be at CERN, and the difficult problem of keeping them actively "involved" in the experiment is widely recognised.

This problem is particularly acute for computing, since this is one area where home-based physicists, with access to a home-based centre, should be able to participate in and contribute to the experiment.

It will of course be true, as is the case in large collaborations today, that most physicists are users rather than writers of the collaboration software. The interface for these physicists is via well defined user subroutines, and they are neither expected nor permitted to tamper with the software "core". Organising and writing this core software will be the task of a small team of physicists who can program, possibly with the help of professional programmers.

This brings us to the thorny problem of how to maintain and run production software at a number of geographically wide-spread computer centres comprising several different makes of main-frame. Anybody who has ever been involved in this kind of software management will stress, vehemently and at length, that any degree of success demands a high level of organisation, and the adherence to well defined standards by those writing main-stream analysis software. What was adequate 10 to 15 years ago, in small groups in which the software expert was just "down the corridor", is **NOT** sufficient today! At the recent software workshop (refs. 2, 3) held at CERN it was interesting to see the care and attention lavished on software management in other fields (e.g. European Space Agency, JET Fusion Project, Britoil). It was also depressing to realise the extent to which much of present day HEP software management languishes behind. However there now exists considerable experience[1] in how to tackle this problem, and there are surely grounds for optimism for the LEP era.

In summary the following points have been stressed repeatedly in WG3 discussions:

- the crucial importance of tackling software design (structure, modularity, interfacing, I/O, different machines) at the start, **before** one line of code is written. From this point of view, the actual calculations and algorithms are trivial.

---

[1] i.e., remembering Oscar Wilde's epigram, most of the worst mistakes have been made, at least once.

- the need for a machine-wide source code maintenance system. This function is presently filled by PATCHY. The conceptual framework of PATCHY is still widely supported; the detailed implementation reflects the card-reader era, and needs quite a major update. A PATCHY upgrade is being discussed.
- documentation is the sine qua non. It is of course an obvious part of software management, and arguably the most important part. One speaker at the recent software workshop (not an HEP physicist) remarked, "If it isn't documented, it isn't software." One could do worse than emblazon that sentence in every HEP office in Europe.

3.4. Consequences for European HEP computing

We shall now itemize what we see as the major needs for European HEP computing in general, and computing at CERN in particular, in order to meet the demands of LEP experiments.

1. The detailed on-line requirements ("computing in the counting-room") fall outside the terms of reference of WG3, but are covered by WG1. We would just remark that physicists are very conscious of the problem of **monitoring** several hundred thousand channels. Even if on-line computers are powerful enough to generate several thousand histograms, the physicists cannot look at them all! Furthermore the appreciation of which are the key histograms comes only as data taking experience is acquired. The **presentation** of on-line information will repay careful study.

2. At any stage along the chain from raw data to DST information, it will be essential for the physicist to be able to "scan" events in order to monitor hardware and software performance. This implies a powerful screen/hardware/software combination; the UA1 Merlin device is the forerunner. The implementation of such an event scanning capability touches the issue of interactive computing which is discussed in detail later in this report.

3. We have estimated that each LEP experiment will need ≈ 4-5 IBM 168 equivalent cpu power with an uncertainty factor of ≈ 2. The present CERN configuration (IBM+CDC) is worth about 9 IBM 168's, but one should subtract at least 2 units to allow for the operating system,

WYLBUR, MILTEN, JES, TSO, CERNET and other uncharged services. It is present CERN policy that only one third of an experiment's computing should be done at the CERN centre, the rest being done "outside". Continued application of this policy suggests then that CERN should provide ≈2 IBM 168's equivalent per LEP experiment. Taking 4 LEP experiments at 2 IBM 168's each, plus a similar amount for non-LEP computing, and adding something for the operating system etc. as above, suggests ≈20 IBM 168's as a suitable figure for CERN's central cpu power for the LEP era. This would be a doubling of cpu power in the next 5 years, which is not out of line with past growth and the availability of increasingly powerful main-frames, (see figure 5, and chapter 9 on costs). It should also be mentioned that, the one-third to two-thirds policy notwithstanding, CERN should expect to support the bulk of the computing load in the first year or so of LEP experiments, as the first experience of real data is absorbed.

How CERN should provide this power is discussed at some length later in this report. The issues of one or two main manufacturers, emulators in an off-line capacity etc. are sometimes political and even emotional. It should be remembered that for the physicist the question of how in detail cpu power is provided is one of almost complete indifference; it matters only that adequate cpu power be there, **and easily accessible.**

4. Crucial to the effective use of "outside" centres is the existence of **RELIABLE** links between the various centres. A physicist at an outside centre who cannot quickly and easily transfer the latest correction code, calibration constants etc. from CERN to his own centre will go away and do something else. We are not concerned here with the continual transfer of hundreds of megabytes (i.e. tapes), but rather with the occasional transfer of up to a few megabytes. Hence it doesn't matter if such a transfer takes even a few minutes; it does matter if one has to spend half a day trying twenty times in order to effect the transfer.

5. The proliferation of inter-computer links, event scanning devices, networks of personal work-stations (discussed in chapter 6) etc. raises the fear of a parallel proliferation of dialogue languages with which the hapless physicist must cope. This is already a well-known problem, and

has been so for many years. It is difficult to quantify, but we suspect that the loss in working efficiency and in use of available resources is appalling. Is it too unrealistic to **plead** for an HEP computing environment in which the physicist can sit down at **ANY** terminal, and work with the **SAME** set of commands for file manipulation, file transfer, editing, and job submission? The answer unfortunately is that such a Utopia **is** too unrealistic! However it is both possible and highly desirable to have at least a "core" editor available everywhere. This desire is illustrated by the popularity of the "mini-WYLBUR" editor available on the PDP 11 machines.

4. THE PRESENT GENERAL SERVICES AND FACILITIES.

The purpose of this chapter is not to enumerate **all** the general services and facilities at present on the CERN site, but rather to demonstrate the scale of the major components today, as a background to the discussion that follows.

4.1. Central processing capacity and usage.

The left side of figure 4 shows the usage of the CDC 7600 per week in 7600 cpu hours over 4 years. Apart from the record pre-conference week in 1979, where short and debug job turn-round were manually restricted in favour of production work, and 149 cpu hours were achieved, the 7600 remains capable of sustaining 125 to 135 cpu hours per week, delivered to users, when loaded. For comparison with the next paragraph, these figures are equivalent to about 400 hours per week of IBM 370/168 time.

The right side of figure 4 shows the same graph for the IBM and IBM-compatible computers in the centre per week over 4 years. The units of the vertical scale are IBM 370/168-3 equivalent hours, since the computers involved changed during the period. In 1979 we moved from a single stand-alone IBM 168, which achieved 100 cpu hours per week on average, to a dual IBM 168 and IBM 3032 configuration, which achieved 220 cpu hours. In the late summer of 1981, the continuously increasing pressure caused total saturation of the IBM side of the computer room. The effect on turn-round

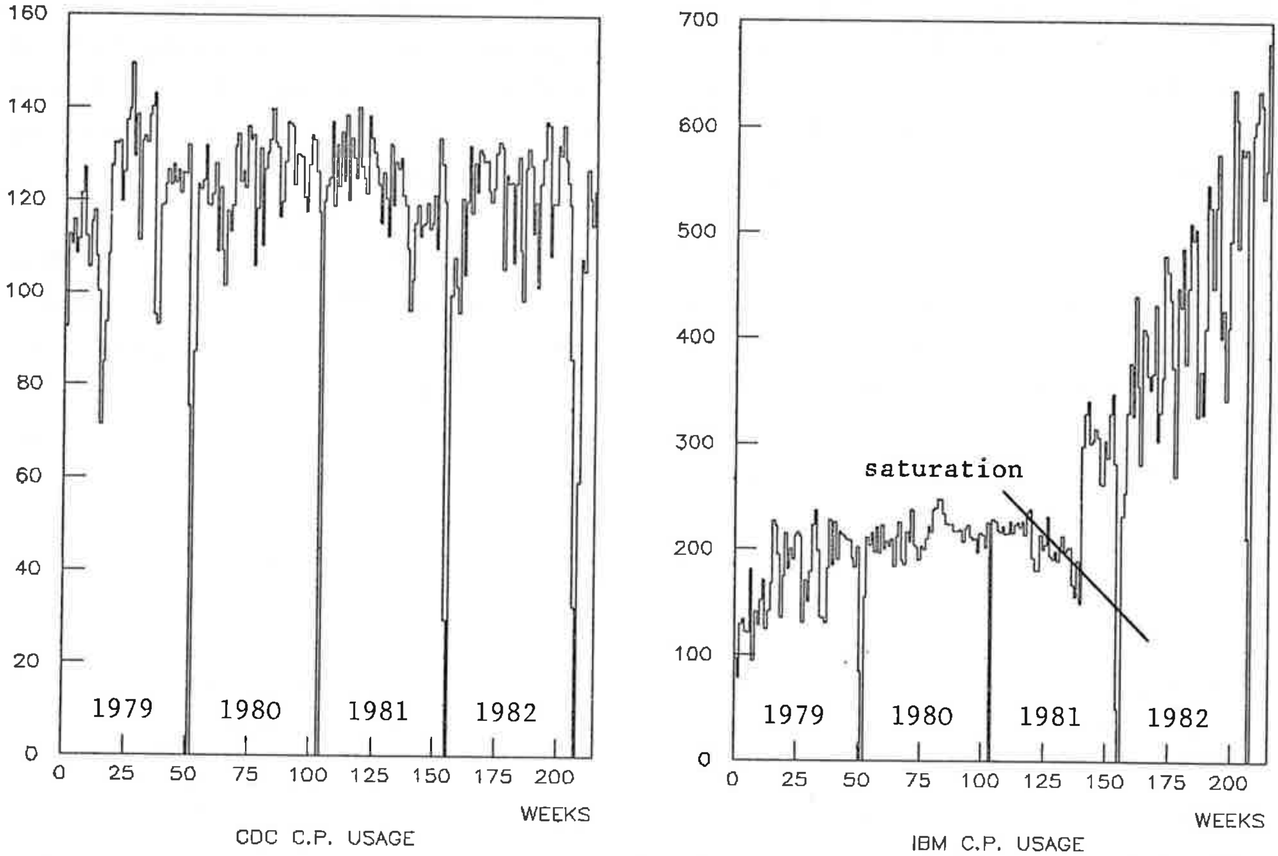

Figure 4: CDC 7600 and "IBM" cpu usage per week, over 4 years.

figures (not shown) was disastrous, but the effect can clearly be seen in the overall cpu figure, where the machines were not able to keep up with the pressure and the production hours delivered to users began to drop sharply.

The installation of the IBM 3081 model D relieved most of the production pressure, but problems remained in keeping up with the demand during prime shift. These were a combination of not having enough disk units and heavy load on the mass storage system, which was also saturated. The acquisition of more disks, together with tuning of the system and the mass store, (improved utilities, and changed methods of working via individual contact with users), has relieved that pressure, but as we will see the demand continues to grow. Both the prime shift and the mass store remain on the limit of saturation. The installation of the Siemens 7880 (Fujitsu M200) in June 1982, together with the upgrade of the 3081 to a model K, have provided sufficient production capacity for the present, but as one can see from the graph, the load is growing fast, if erratically. On the basis of the graph we shall again be saturated towards the end of next year, although the proposed modernisation of the CDC equipment could increase the CDC part of our capacity by about 50%.

Figure 5 gives the total cpu usage of central computers at CERN since 1958 in units of 7600 user hours per year. The usage in 1980 to 1982 is clearly below the straight line,[2] but heading back towards it,[3] and the saturation limit of the present installed capacity is indicated. The extrapolation of the line to 1987 gives a figure approximately two times higher than the WG3 "minimum need" number of twice the present installed capacity. This is in line with all the comments we have received from physicists, in letters, open meetings, etc., be they in LEP experiments or not. No-one at all thinks that our figures are an **over-estimate**. One further corollary is that capacity **must** be installed in time to meet the load. The negative effects of saturation and the immense amount of human time utterly wasted, by users and computer centre staff, must be avoided at all costs.

Figure 6 gives the distributions of cpu usage on CDC and IBM by CERN Division during November 1982. The two distributions are very similar, with the exception of the greater usage of CDC by the Theory Division, and show more than 93% of the cpu being used for physics. (The slice labelled operational support on the IBM side of the graph represents the cpu time used to perform file base backup and other general purpose service facilities).

---

[2] The line is actually a fit to the years 1968/1980, since the graph was first produced in 1981.

[3] The 1983 point is an extrapolation based on the actual usage in the first two months.

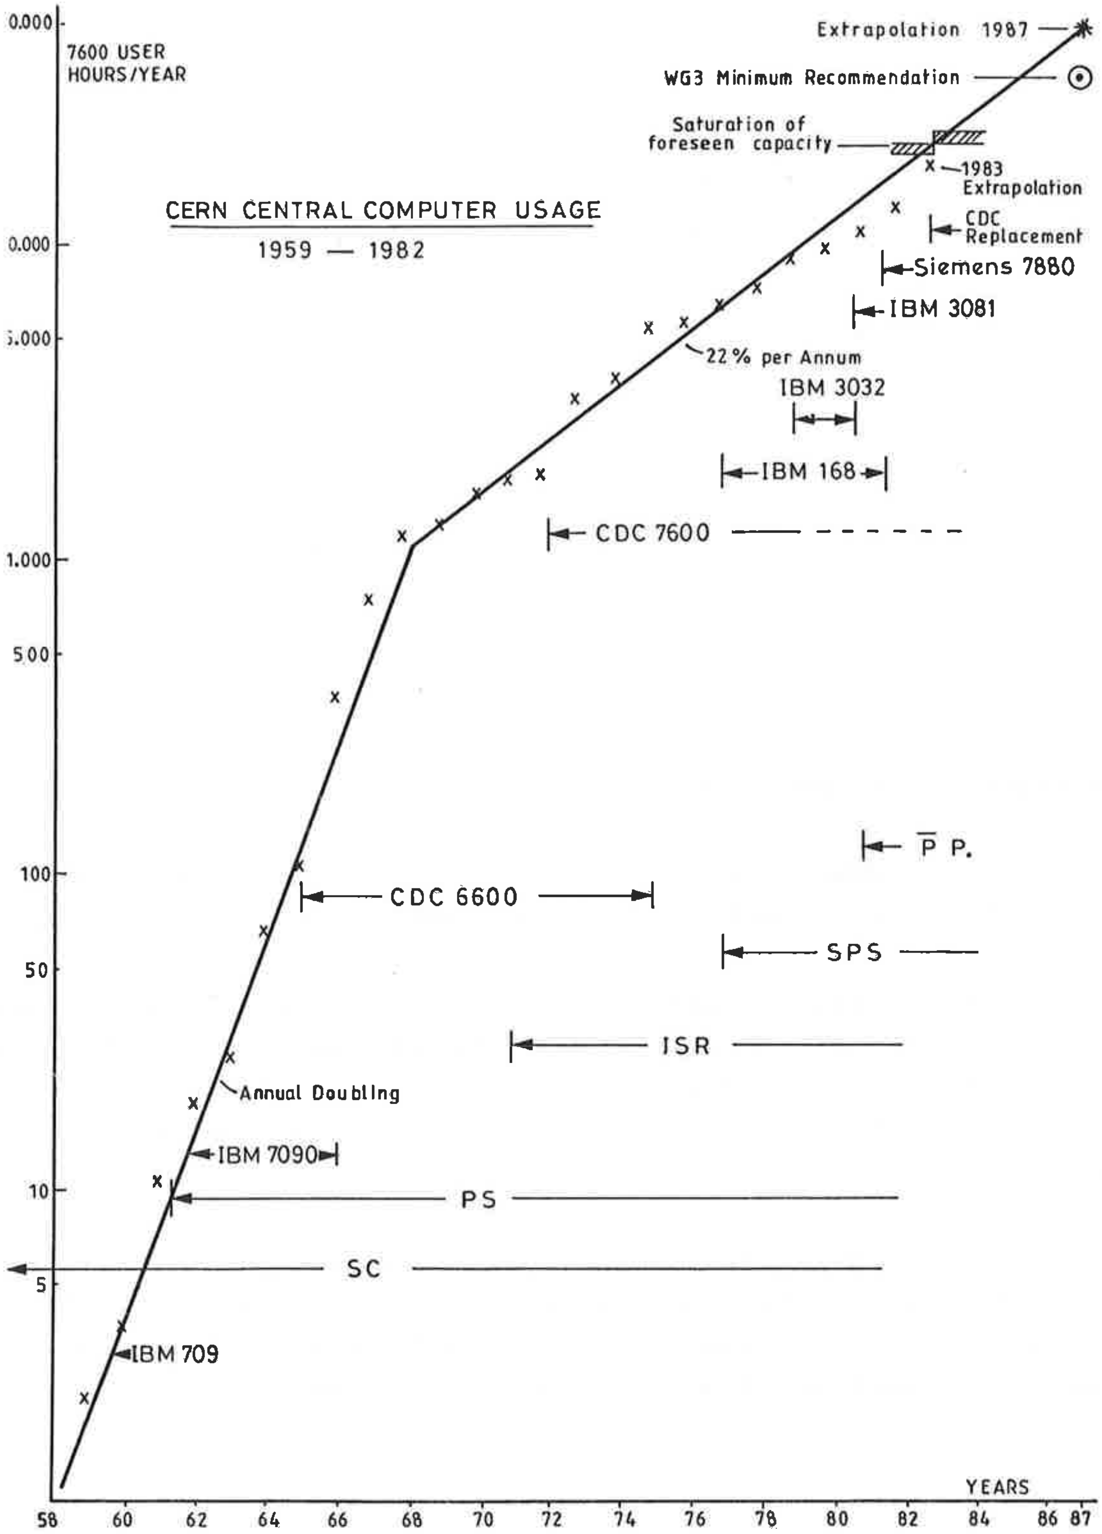

Figure 5: Total Cpu usage at CERN since 1958

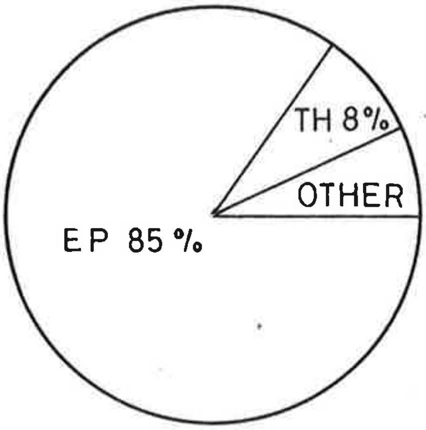

7600 CPU TIME BY DIVISION

IBM CPU TIME BY DIVISION

Figure 6: Distribution of cpu used by CERN divisions on CDC and IBM

4.2. Numbers of Jobs and Users.

As a further indication of the scale of the present central operation we give some graphs concerning numbers of jobs and numbers of users.

The left side of figure 7 gives the number of jobs run per week on the CDC 7600 over 4 years, specifically excluding the jobs run on the CDC front-end computers. It shows a small decrease over 4 years, presumably from a minor migration of small jobs to the IBM side, but the 7600 still runs about 12 500 jobs per week.

The corresponding graph, (shown on the right side of figure 7), for IBM[4] shows a rather spectacular rise, and is a witness to the constantly growing pressure which requires regular tuning and installation of new equipment, such as teleprocessing lines and disks, in order to keep pace.

---

[4] The definition of a job changed during the period, when list offline jobs ceased to run as batch jobs. Nonetheless, for consistency, list offlines are included as jobs all through this graph.

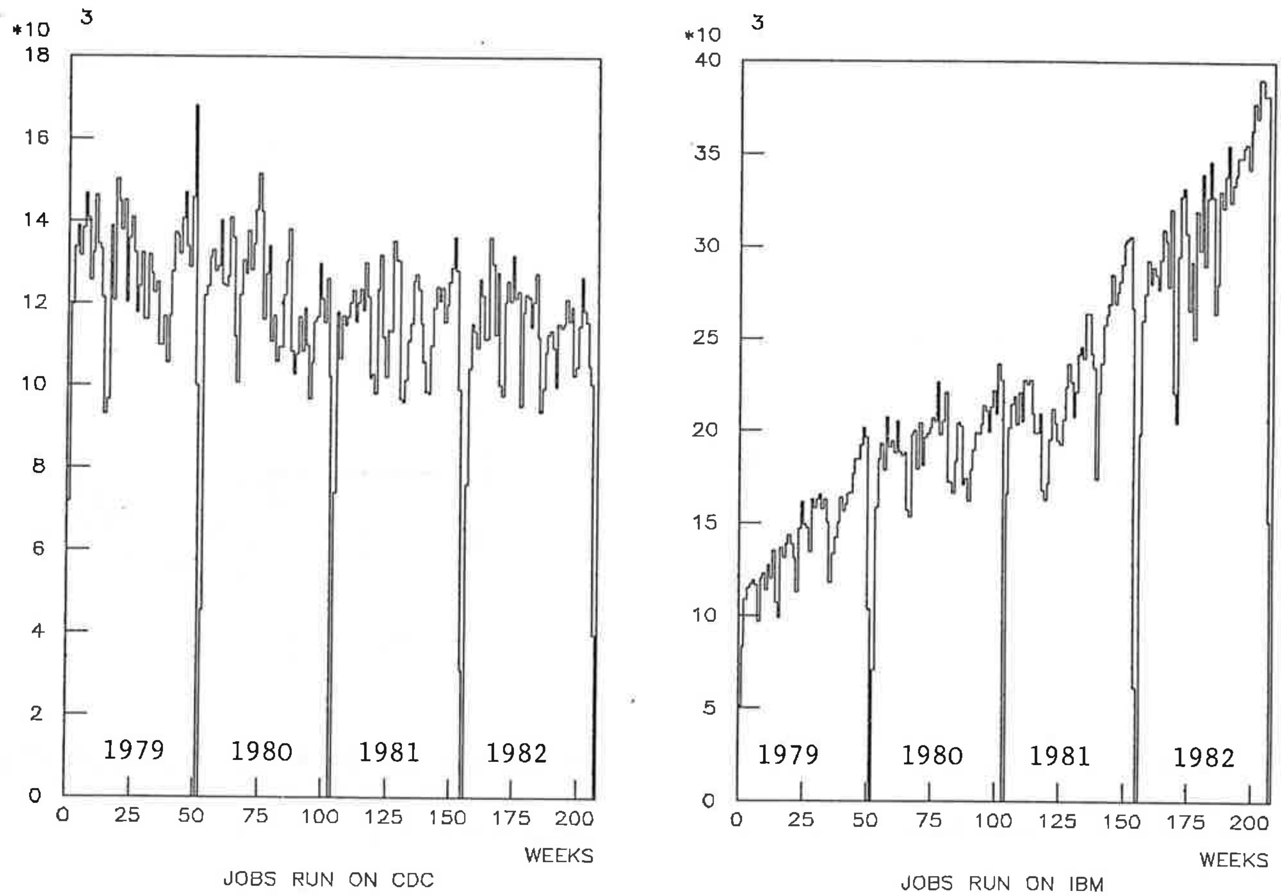

Figure 7: Numbers of CDC7600 and "IBM" jobs per week, over 4 years.

In order to give some idea of the distribution of the workload, we give, in figure 8, the distributions of cpu time and numbers of jobs by job class on the IBM system during November 1982. The major job classes are as follows (in 168 equivalent time units):

- Class X: Express, no tape, up to 15 secs.
- Class S: Short, no tape, up to 2 mins.
- Class T: Short, tape, up to 2 mins.
- Class M: Medium, up to 15 mins.
- Class L: Long, up to 2 hours.
- Class V: Half-price low priority, and jobs between 2 and 4 hours.

They show that over half the jobs are in the express category, using less than 2% of the cpu, although they use much more in other system resources. At the other end of the spectrum, the long and very long jobs consume 67% of the cpu whilst contributing less than 5% to the number of jobs.

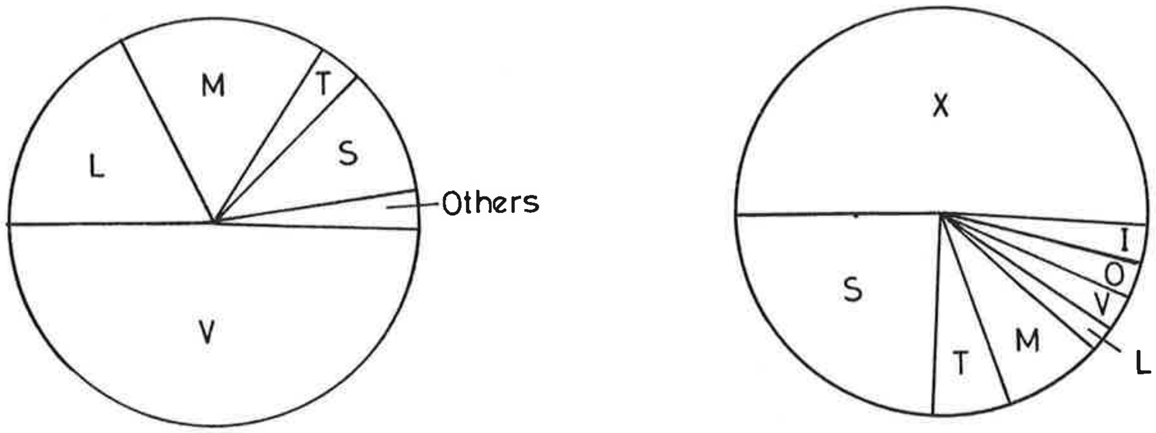

IBM CPU TIME BY CLASS   NUMBER OF IBM JOBS BY JOB CLASS

Figure 8: Distributions of cpu time and numbers of jobs by job class.

The numbers of users per week running batch jobs on the 7600 and "IBM" systems over the past 4 years are shown in figure 9. Whilst the 7600 users have decreased slightly in a manner similar to the numbers of jobs per week, the number of IBM users has risen sharply. The summer plateau and the autumn increase of new users are clearly visible, and the latter contributes regularly to tremendous pressure at the end of the year on support staff. In addition we were obliged again this year to install more WYLBUR lines to reduce the queues for access during prime shift.

Finally, the distributions of users by division using either batch or time-sharing facilities, on both IBM and CDC, (not shown), demonstrate that the non-physics divisions constitute a much greater proportion of the users than their cpu time usage would indicate.

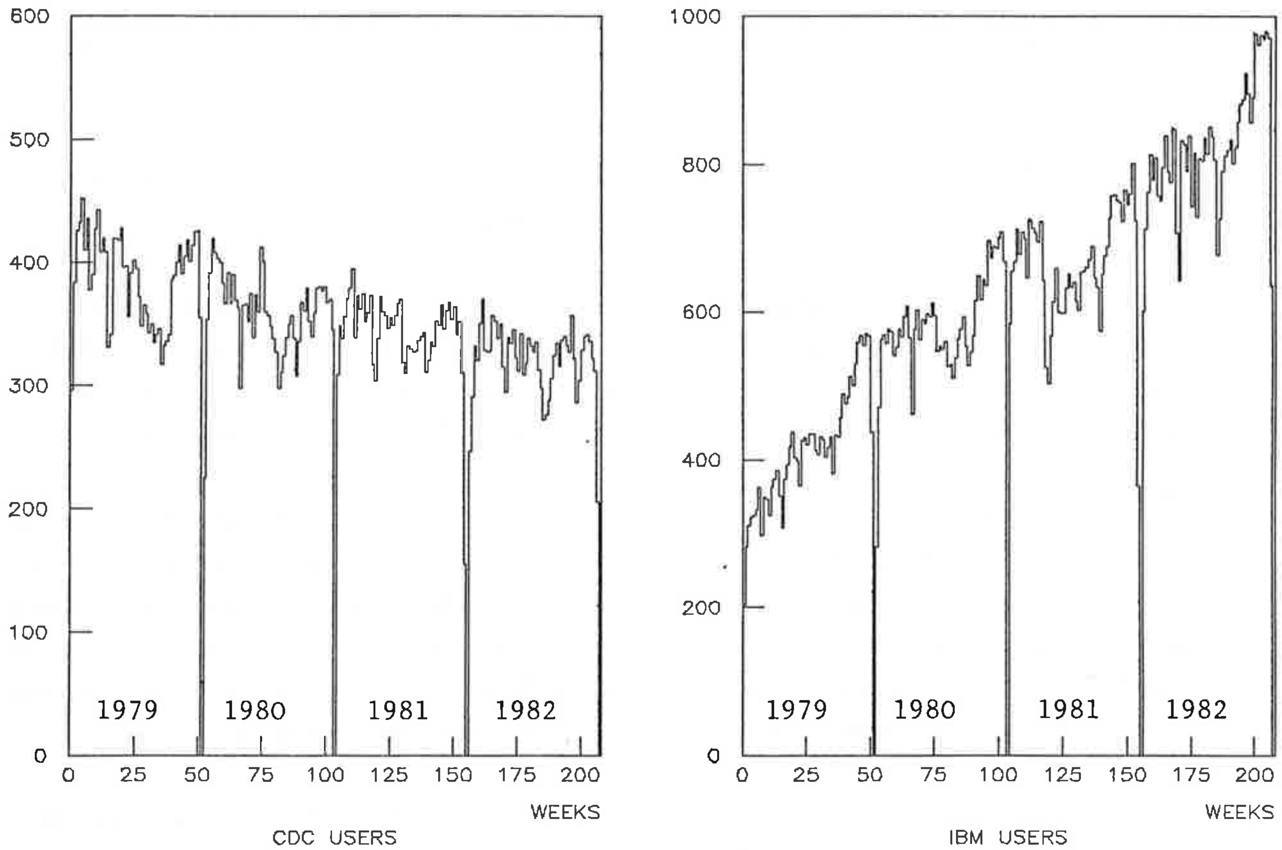

Figure 9: Numbers of users running batch jobs on the 7600 and on the "IBM" per week, over 4 years.

4.3. Summary of Computer Services offered.

The purpose of this section is to give a brief list of the major computer services offered today as background to the discussions that follow in the report.

1. The CDC 7600 offers powerful batch processing, excellent debug turn-round,[5] fast floating point arithmetic, precision, and wide range of

---

[5] It is a particular policy of the CERN computer centre to provide the best possible turn-round for express and short jobs in order to save human time. To this end great attention is paid to the scheduling on both the 7600 and the IBM system.

numbers. Its major weakness today is its small memory, and many modern experiments have difficulty, or are unable, to use it for production.

2. The "IBM" system offers powerful batch processing, excellent debug turn-round, virtual memory, a wide range of applications, the latest peripherals such as mass stores and laser printers. Most new devices are more readily available in a form where they can be interfaced to the IBM than to other computers. Its problems are its antiquated job control language, its difficult error messages and utilities, the care necessary to avoid precision problems, and in some cases the range of the numbers.

3. The main job of the CDC front-ends is to feed the 7600 and the remote I/O stations. They also run the INTERCOM time-sharing service. The latter is fully interactive, but lacks a real EXEC facility, which decreases its usefulness. It is limited in memory, facilities, and the number of simultaneous users it can handle. Nonetheless the front-ends run many special applications which are hard to move elsewhere.

4. The WYLBUR time-sharing service, running on the IBM system, provides a very popular and cost effective service for 180 simultaneous users. Over 1450 accounts were used per week during February 1983 on WYLBUR. WYLBUR is **not** fully interactive, (you cannot communicate with a running FORTRAN program), but it provides an extremely versatile EXEC language, which is widely exploited. The limits of WYLBUR are beginning to show, e.g. the line-oriented editor, (nonetheless one of the best of its kind), and the excessive use of WYLBUR's inefficient EXEC files in applications for which they were never conceived.

5. We offer a form of fully interactive TSO service on the IBM, via a bridge to WYLBUR. Fears that TSO is a heavy resource-consuming system have led us to limit this to physics applications such as event scanning, physics databases, and the graphics system HTV. As offered, the service provides editing, and file manipulation etc. via WYLBUR, whilst the interactive program itself is launched under TSO.

6. The full TSO interactive system, including editing and file maintenance, is provided to systems programmers only, for whom it is essential because of its integration into the heart of the system.

7. The IBM mass storage system, MSS, and its associated software for managing files, HSM, is an essential service provided to all computers on the site which are connected to CERNET. It allows smaller disk configurations to be installed at the minis than if we did not have an MSS, and the safe backup by the computer centre staff is a major service. Its effect in containing the tape mounting load has been significant.

8. Other special services. The centre runs a number of special services such as the VAX computers for the LEP database system and the LEP Computer Aided Design system.

4.4. Human services.

Lastly one must not forget the human services provided by the computer centre staff, such as:

1. Operational Services. All tape services. Manning Consoles. File-base backups. Printer, microfiche, plots, delivery service, etc.

2. Systems Group. Installation and maintenance of systems. Support of local systems such as WYLBUR. Local modifications to standard systems (essential). Tuning for the CERN workload. Diagnosis of problems.

3. User Support Group. Program Enquiry Office consultancy. CERN Program Libraries on CDC, IBM, VAX, APOLLO. Documentation. Standard procedures, utilities, and EXEC files. Accounting, registration, control of budgets for money and space. Documentation.

Rather than explain at length the necessity and value of these services, let it suffice to point out that **none of these services can be provided by the manufacturers.**

5. OPTIONS FOR THE NEXT DECADE

5.1. Introduction

The purpose of this chapter is to present some of the discussion that has led to our conclusions. Much of it has already appeared in the Interim Report of WG3, (WG3/cmj/82/30, revisions 3 or 4), but a number of arguments have evolved since that time. In particular, the very crucial question of how to provide for Interactive Computing has been placed largely in a separate chapter, as befits its importance.

5.2. Distributed Computing versus the Computer Centre

There is a popular idea that distributed computing will replace computer centres in the next decade. Being cynical, one could observe that this popularity is inversely proportional to the success of the centre best known to the reader. In particular, there is a trend that computer scientists find their own overloaded university computer centre to be unacceptable, and are able to work quite happily on a small private system. They rarely need large central processor power. Hence, the idea that distributed computing may replace computer centres entirely may not be wrong in some applications during the next decade, because of the power of even the present microprocessor systems. However, it is likely that high energy physics is one of the applications where such a process will take a very long time, just because of the vast amounts of central processor time needed. Furthermore there are strong reasons for central siting and operation, such as efficient utilisation of processors, operator effort for mounting tapes, economy of scale, expensive peripherals such as mass stores, bulk laser printers and large plotters (see WG3/jmf/lg/82/9).

Hence, the conclusion of WG3 is that, in the next decade, we shall have neither just distributed computing, nor just a computer centre, **we shall have both.** We shall return later to distributed computing, and concentrate for the moment on batch processing in the computer centre.

5.3. The Computer Centre

Our discussions have suggested a factor 2 increase in central computing power in time for the start of LEP experiments (1987), but, as before when CERN has made this calculation, there are large uncertainty factors in the load and need. No one in the open meetings or the working group would support less than a factor 2 increase for batch processing, and many thought more.

5.3.1. IBM Compatible Computers

The question of which computer should provide these central cpu cycles is nowadays far less emotional than, for example, at the time when the CDC 7600 system was purchased. At that time the full system cost of the order of 40 MSF (in 1971 francs). The recent acquisition of the IBM 3081 Model K system of slightly greater power cost of the order of a quarter of that price in francs (depending on what peripherals you wish to include in the cost), which is less than a quarter in terms of real money. Our option to add a second processor to the Siemens-Fujitsu machine would add 2/3 of a 7600 for less that 1 MSF. (Beware that this is a very special case). A further example would be the doubling of the 3081 K, to the 4 processor 3084 Q, for about 7.5 MSF. Note also that increasing the power, requires increasing the number of disks to maintain a balanced configuration. This was made abundantly clear when we installed the 3081 before the money for disks was available.

Nonetheless, on the IBM-compatible side of the room, the planning can be relatively simple. Raw capacity for production can be added incrementally when needed at incremental cost with no problems of user conversion. This is expounded in the costs and manpower chapter. Furthermore, the opportunities for competitive tender for IBM-compatible equipment are promising, with several European firms available, albeit selling Japanese equipment. IBM itself cannot be ignored, especially as the 3081 was manufactured paradoxically entirely in Europe, even though many label it as American. The support that IBM provides, necessary for the software, firmware and specific peripherals, should not be overlooked, and this in itself means that it could be very unwise to become too heavily just "plug-compatible". A good balance is necessary if CERN does not wish to find itself in a situation where greatly increased software support is demanded of CERN manpower.

A further factor in this situation is the existence of emulators like the 168E, the 3081E development which will be easier to use, and the 370E. The bank of 5 168E's connected to the centre via CERNET has produced significant capacity[6] this year. The working group is unable to agree whether 3081Es should play a significant off-line role, the true cost in manpower, money, and central computer resources being hard to express. It would be fair to say that the case for emulators in off-line production is not proven either way at this stage. Some members of WG3 are strongly against this off-line use, feeling that the emulators' role should be on-line, and that in the time taken to develop a new version the cost benefits become marginal. Others are not sure. It probably needs another round (the 3081Es) for the case to evolve. The real point is that here also some processing power could be added incrementally as needed up to a point.

5.3.2. CDC Computers

The particular case of CDC is more difficult and involves many factors, some of which are political. The 7600 has been at CERN more than 10 years and is still, for many programs, our fastest cpu. By now it is expensive to maintain and consideration of its replacement is very topical, since we are in the process of negotiating a contract for new machines. The problem has been that CDC decided some while ago to drop the 7600 and its operating system, and there is no obvious successor, i.e. a bigger machine with a compatible operating system and a technology more modern than transistors. In fact, leaving the vector processing CYBER 205 computers aside, CDC has no cpu today faster than the 7600. Indeed it will still offer CDC 6000 series architecture, built of 7600 transistor modules, for the top end of their range for the next couple of years.

Many users still appreciate the excellent 7600 service for fast debugging of short and medium jobs, and its precision and range of numbers are important for some problems. The major difficulty for the batch production of physics events is that today's experiments can no longer fit their programs in the limited 7600 memory, and even those that can fit in lose speed swapping

[6] The capacity of the 5 168E's is of the order of 2 * 370/168. In practice the three physics groups using them have achieved a production equivalent to about 3/4 of that attainable in the same time on one real 168, which is less than half of their potential full capacity.

overlays and superlays. CDC does not offer a virtual memory system today, and the 60-bit word length is an irritant now that 8-bit bytes are standard on most computers and data acquisition systems.

The outside Universities and Laboratories that bought CDC, largely to follow CERN, have also to be considered. However, since they run 3 different CDC operating systems outside, whichever we choose, we can only be compatible with one of them. At least the FORTRAN compilers would be the same. The internal CERN CDC users are divided broadly into two sets. One set, the physicists, uses mainly the 7600 for their data analysis programs, which in large collaborations must necessarily run on a broad spectrum of machines, including IBM. Conversion to IBM would not be a problem for most of them. The other set of non-physicist users tends to run only on the front-end CDC machines with many small interactive applications. Conversion for them would take time and they would need help. However it is true to say that the demand for interaction is such that many physicist interactive applications, especially event scanning, wanted to use the front-end machines, but were defeated by memory requirements well above that available.

The question has been whether to reduce our CDC presence to a rather small machine for program compatibility testing and the existing interactive front-end applications, or to invest in a new generation of CDC machines of comparable or larger power. We clearly would like a machine of at least the power of the 7600, with virtual memory and the 64-bit architecture that CDC is suggesting they could release in a couple of years. This is just not available today. Hence, we have negotiated a stop-gap solution which will involve transistor machines as discussed above, using NOS/BE (as on our present front-end machines, and involving conversion of the 7600 users). The alternative system NOS, (which we have never run), would involve a much larger conversion problem, and has thus been ruled out. The plan is to introduce a CYBER 875, of the speed of a 7600, but having no front-end, around July 1983. It will be accompanied by a small machine, a CYBER 835, which will not function as a front-end, but will provide some back-up, access to files when the 875 is unavailable, and interactive INTERCOM computing at about the present level. The users of the 7600 will have to be converted to the NOS/BE system on the time-scale of Christmas 1983, when the 7600, 720 and 730 will disappear. At this point the 875 will be upgraded to a double processor.

The first benchmarks we have run show that the 875 does indeed run at the speed of the 7600, and that the memory problems will be definitely less. There is certainly a category of problems, that run badly on the 7600, because of the time spent in swapping overlays, that will run well on the 875 because its gives a real memory increase of the order of a factor 2. Nonetheless, we still see problems in being able to run the very latest generation of experiment processing codes such as UA1 because of memory size. This will still be true even with a number of local modifications to the system to further minimise the memory problems. These modifications are essential in order to exploit the machine at a reasonable level.

Taking into account the architecture, memory size and operating system of the 875 dual-processor, as well as the fact that it will have no front-end computer, we estimate that it will produce 1.5 7600 equivalent units of power. The net gain to CERN will be 0.5 7600 units, which are equivalent to about 1.5/1.7 IBM 168 units on average. The real decision will come in a few years time, when we may have the option of converting to a CDC 64-bit based operating system, with virtual memory.

5.3.3. Vector Processing Machines

WG3 has been unable to spend much time considering vector processing machines, but we can make the following points.

The CERN benchmark, consisting largely of a set of FORTRAN high energy physics data processing programs, was run two years ago on a CRAY 1, and repeated during October 1982 on CRAY 1M, 1S and X-MP computers. We confirmed that, for these programs, the 1M and 1S machines run only at a speed of about twice the CDC 7600, and that the X-MP runs a little faster through its improved cycle time (9.5 instead of 12 nsec). Since a factor 2 is close to the cycle time ratio of CRAY 1/7600, it is clear that advantage is not being taken of the vector facilities with these codes as they are written today.

The paper of Zacharov (ref. 1) suggests techniques (albeit involving major investment) that could be applied to the above processing codes in order to exploit the vector capabilities of such machines, but it is fair to say that, up to now, no physicists have applied such techniques to these problems. They have preferred, where any effort at parallelism has been attempted, to attack the seemingly easier problem, which requires only minor reorganisation of the program, of using multiple 168E emulators. Hence we do not know how successfully such machines could be used for physics data processing.

The CRAY produced more spectacular results for a CERN accelerator design program, which achieved a factor of 10 over the 7600, and for some important QCD lattice calculations which we are currently running. Here, an expert making a short-term investment has produced a factor 25 over the 7600 for the central part of the problem. Furthermore, the successful QCD lattice example that was run during the nights and weekends over a couple of weeks on a CRAY is but one case amongst many that it is desirable to run. To attack this problem more seriously will require a very large number of hours indeed, even on the CRAY. It is out of the question on our fastest scalar computer.

Alternative approaches to these QCD calculations, in other institutes, use either array processors, (e.g. Floating Point Systems), attached to channels on scalar computers, or propose to build multi-microprocessor home-made systems. These solutions, in our opinion, can never compete with a vector processor computer, because the problem is not just one of central processing power, but also has major requirements of memory and data flow. In the case of the add-on array processor, one must write and read the data along the channel to the processor before it can do any work, and the start-up time is very considerable. The multi-micro solution requires that each micro is equipped with a very considerable amount of memory for each processor, which then unbalances the costing of the otherwise cheap microprocessors involved. Furthermore the work involved to build such engines will be considerable, and not achieved in a short time. The CRAY solution works today.

Thus there is a small and growing number of problems in high energy physics today that justify, and indeed require, a vector processing machine for their solution. We believe that there will be more in the future. Hence the problem today is how to get a foot in the door and gain access to a CRAY, (or possibly a CYBER 205). We see a very real possibility that, soon, it may be perfectly reasonable that one of the high energy physics laboratories in Europe purchases such a machine for all the European physicisists who could use it to solve their array-oriented problems. This may or may not be CERN itself, but clearly CERN is an obvious candidate.

Another point concerns a new and coming generation of scalar computers that include array processing instructions in the cpu, as opposed to those on the end of a channel. This exists today from Hitachi, and we have strong reason to suspect the arrival of new offerings from other major manufacturers. This could well be combined with a development in software, for example the addition to FORTRAN 77 of the proposed FORTRAN 8X facilities for array processing. This would provide a natural route for physicists to bring array processing into their codes, and we should keep this option under particular study.

A final point, of which we are not sure, concerns the Monte-Carlo calculations for LEP physics. We have seen that, in general, the fraction of the total computing taken by Monte-Carlo has been much greater for PETRA experiments than for most of the PS, ISR, and SPS physics. At our open meetings some physicists have suggested that for LEP, Monte-Carlo could be even much greater than the data processing. We suspect, without being sure, that some of this Monte-Carlo lends itself more easily than the event processing to the exploitation of vector machines, perhaps using the techniques suggested by Zacharov (ref. 1).

5.3.4. Major Peripherals

There are good arguments for the central siting of most expensive peripheral equipment, provided of course, that these are easily accessible to all client computers and users on the future networking system. Obvious examples of fast laser printers, microfiche (or film?) devices, and large plotters have been given already. One must not forget either the time lag between the availability of new peripherals such as laser printers, 6250 bpi tapes and probably video disks on IBM systems and their arrival on minis.

Given that a large number of distributed computers exist now on site, and that this number will increase, the really key item is the Mass Storage System. We must foresee either a substantial expansion of the existing IBM MSS (it could be increased by nearly a factor 5), or the purchase of another mass store during the next decade if such a device in a more desirable technology became available. It is also clear that we must envisage a substantial effort in software to achieve a highly desirable network-wide filing system, which is far more than the existing ad-hoc mini-computer back-up schemes of today. The connection with the ERASME VAX today is an example of such a system. However for a general implementation a great deal of redesign would probably be necessary.

The demands from multiple machines on the site imply that we should think less of an IBM attached mass store, and more in terms of a Mass Store as a general "File Server" as well as a fast connection between computer centre machines such as that available from Hyperchannel. This implies, however, purchasing a new IBM compatible machine to run that file server. The justification comes from the large number of computers on the site that use the mass store as a repository for their files. These files must pass through the central IBM computers on their way to the target machines, even if these files have nothing to do with the IBM itself. The result is that a significant fraction of the cpu cycles of our large processing machines is spent in the overhead of transferring those files from the mass store to CERNET. Furthermore, the throughput of this process is limited by the priority that can be assigned to the system tasks in question, given that there are other high priority tasks in the machines such as WYLBUR. It is therefore logical to consider whether we should install a smaller machine which is dedicated to running the file server task at highest priority. This needs a study of the trade off between wasting the expensive cycles of the the big machines and the cost of a dedicated smaller machine. The study must consider how this traffic will evolve in the future, for example now that the minis on site can afford to buy disks much larger than in the past and hence generate potentially much greater back-up traffic. We already know that the present transfer throughput is insufficient to avoid inconvenient queues during the day.

WG3 has not had the time necessary to study these particular problems in sufficient detail. However we recognise that they are important, and will become even more so in the future. We recommend that two associated studies be made. One should consider the role of the mass store as a central file server for the whole network, including a network-wide filing system, and the justification for a separate computer attached to the mass store to avoid wasting the cpu cycles of the major processing machines. The other should consider a high speed connection between the many computers in the centre. This should compare the present use of CERNET, the commercial offerings such as Hyperchannel and LCN, and the possible future CERN spine network.

We have explained that, when installing new cpu capacity, it is necessary to install new peripherals as well, in order to realise that capacity. The detailed list of peripherals that are required, together with their justifications and pricing, are presented in the chapter on costs and manpower.

5.4. Distributed Computing

In the working group paper WG3/mgnh/82/11, M.Hine has pointed out the need to distinguish between 3 types of distributed computing which have different consequences, e.g. on the centrally organised support staff. In practice one needs to separate:

- Physically distributed computers
- Logically distributed computers, and
- Managerially distributed computers.

The CERNET node computers are physically and logically distributed. They are centrally managed. They provide an integrated set of end-user services. The real dangers come from the last category above. The demands to interconnect and integrate many managerially distinct services can be a potential "black hole" for support manpower as we explain below.

The expansion of distributed computing is inevitable, and can bring tremendous gains if done well. The dangers however are many and real. The integration of all these distributed, and different systems, that will be purchased from many different budgets, can be a very difficult task. It demands work and knowledge that cannot be purchased from the manufacturers. In an extreme case, one might think that by going just to one manufacturer for everything, one might be able to buy an integrated system, but this scenario is absolutely excluded at CERN, and we must expect very many manufacturers' equipment, of many diverse functions to arrive on site. Central control of acquisitions (even small), of a greatly increased nature from that at present in force, could limit some of the dangers. WG3 believes that this is a very important point.

Early choices of CERN standards can also have a limiting influence, and are hence essential. However, we cannot have absolute control over the visiting groups, we can only show them the value of standards. Even then policies in home countries (e.g. the British support for PERQ) could over-ride standards in some cases. We therefore must be prepared to integrate all of this equipment, and it will be a great deal of work. **It is the integration of the end-user services on distributed systems that can make all the difference between a successful cooperative network of equipment and total chaos.**

5.4.1. The Human Interface, A New Component

The major new component in computing for the next decade is the development of the human interface, something that has been very neglected in the first decades of computing. In the past, systems have been built for specialists only, with unnatural command languages and ridiculous error messages. They have been hard to learn and to remember. Few people think of using the terminal as a natural part of their life in the same way as they use the telephone. Indeed, the first 30 years of the history of the telephone make an interesting comparison with the early years of computing, in that people took that long to understand really how to use the telephone as it is used today. A similar comparison can be made with the invention of the printing press, which was at first thought to be a way of eliminating copying errors. It took 50 years before man realised that the printing press could produce **cheap** books in volume, and that produced a revolution in thinking.

Hence, it is probable that we have not yet fully understood how to use computers, and it will take at least another decade before we do. We believe that the greatest change in the next decade will be in the human use and human interface of computers, and that this will be a new element in the cost and manpower budgets for which we must provide. However this is not to say that we will pay more and get nothing in return. There will be savings because end-users will become more efficient thus either saving manpower or achieving more, and computing will spread its benefits into many more activities than now.

One thing is sure: the interface will provide much more immediate and powerful **interaction** with the running program or application, with correspondingly better presentation of results to the user, graphically on screens and on paper. So far interactive computing at CERN has been limited to a few types of problem and of computer, and has been expensive and, often, usable only by specialists. Our options as to how to provide interaction are expanded in the next chapter.

5.4.2. Multiplicity of Choices and Environments

At this stage it is appropriate to emphasise a key point. A major directive to WG3, to be found in the minutes of an early meeting of the LEP Computing Planning Group, was the need to seek for unification in order that people do not have to learn many different systems. We were specifically requested to find fewer systems than today, preferably just one. This is not new. At a User Meeting around 1970, someone received a round of applause for the remark:

- "the average user can remember only 1.1 operating systems!"

This remark received great sympathy at our open meetings, and it has come up in meetings of the other working groups as well.

Whilst WG3, and indeed all the members of the existing computer centre, understand and support this as fully as possible, it should be clear from what has been explained so far in the section on distributed computing, and from the text of the next chapter on Interactive Computing, that the trend in the World is in the opposite direction. We are faced with a multiplicity of options, and each has good reasons to support it. At our open meetings there were supporters for every choice, and no one was willing to give up his preferred system.

Up to now there has been a mechanism for unification, in that one could only justify one major installation and that was used by everyone. Hence the central system was **the** system. However this meant a choice had to be made of the best single user environment to do all things for all men. And some were better than others... And none of them did everything. By and large WYLBUR was an excellent choice at the time. It has great depth and flexibility and the applications that have been written in WYLBUR execs are a tribute to human enterprise. Nonetheless, everyone can see examples where specific problems are handled much better on specific systems tailored to the job. In summary, one has to accept that the day of the one great super-system that does everything for all people are over. So what do we do...?

It is correct to say that this, in a nutshell, is largely why we have not decided earlier what to do about interactive computing. We have just not been convinced about the superiority of any of the solutions for the general problem. For this reason, and because we believe that interactive computing is of fundamental importance in improving the human interface in the next decade, we have devoted the whole of the next chapter to it.

6. HOW TO PROVIDE FOR INTERACTIVE COMPUTING?

6.1. Introduction

WG3 has spent a substantial part of its time on the choices for interactive computing, as can be seen from the working papers. In particular, this chapter makes use of material contained in the papers on models for the future by Brun, (WG3/rb/82/19), and Jones, (WG3/cmj/82/25). It also draws on the WG3 Interim Report, (WG3/cmj/82/30), and the analysis of Robertson, (WG3/lmr/82/45). In the interest of keeping this chapter reasonably short, many of the arguments of the above papers are omitted, or presented in abbreviated form. The interested reader is invited to obtain the original papers.

6.2. The need and justification for interactive computing

WG3 considers that interactive computing can make a very significant contribution to the important problem of improving the human interface. The particular characteristic of such an "interactive" system, beyond the fact that it is accessed from a terminal, is that the user can program it himself, and can select programs to run on it which he has obtained from a wide variety of sources. Examples are HTV, event scanning, interactive debugging etc. This is in contrast to "conversational" systems like WYLBUR which are programmable only by specially designated personnel, and where the end user is presented with a limited set of facilities chosen in some centralised way. Such conversational systems grew up largely for reasons of efficiency and cost, arguments which have lost much of their validity with recent improvements in the price and performance of computing hardware.

Up to now CERN has provided rather little "public" interactive computing, except a low-level INTERCOM service on the front-end CDC machines (which limited the applications in power and even more in memory), and a TSO service linked to WYLBUR which is restricted to physicists needing interactive graphics to scan events etc. On the other hand interaction is a well accepted feature of specialist systems like Merlin, of accelerator control systems, and of course the data acquisition systems (the numerical majority on site). We have not provided a fully supported interactive system for "public" use largely for fear of the cost, effort, and justification of the use of the central batch machines for such work.

Furthermore WYLBUR, however much it can be criticised in some areas, has been a very successful and cost-effective way of providing many essential semi-interactive facilities. WYLBUR is a very clever system in that it understands what people need to do most of the time and does it so well that many people have not felt the real need for interaction. This has been supported by a strong policy of excellent turnround for short debug work on both IBM and CDC systems, which has been carefully studied and constantly worked upon in order to maintain turnround. Nonetheless we can now see that this form of service has got us approximately half a generation behind many outside computer centres which have launched into central full time-sharing systems, e.g VM/CMS at SLAC and at Rutherford, NOS at Fermilab, TSO at DESY. One should not forget however the effort from the support staff, and the inconvenience of user conversion, that has taken place in some of these sites in order to install such services.

We should not ignore either the tremendous ingenuity going today into by-passing this problem, and the thorough undesirability, (in terms of system resources used), of many of the WYLBUR EXEC files that are being written. These EXEC files are capable of handling a certain range of interactive type problems, but the WYLBUR implementation is inefficient, having not been designed for usage on the present scale. This high usage demonstrates the users' need for interaction. We have many examples where the computer centre has decided that, although certain facilities are needed, EXEC files are not the correct way to implement them. The consequence is that many user groups have then written private EXEC file implementations with probably greater inefficiency than the rejected central versions, and certainly with great duplication. The services provided by these EXEC files could be implemented in a much more efficient manner on most fully interactive systems. Furthermore, being increasingly non-standard, WYLBUR is becoming a growing maintenance problem, especially in the areas of its interfaces with the rest of the system.

The principal justification used for the introduction of an interactive system is the improvement expected in the productivity of its users. While this usually implies a cost benefit to the organisation concerned due to decreased staffing levels, the more important factor at CERN today is the need to improve the performance of the individual user in a period of growing workload and diminishing staff numbers.

6.3. Application Areas at CERN for Interactive Systems

In the section below we list some of the possible application areas for interactive computing, as an illustration of the range of the problem, and of the different components needed in each case.

- Offline Physics Program Development: The basic environment consists of an editor, source management system, Fortran compiler and debugger, and some degree of graphics support together with all sorts of file management utilities. Ideally the environment will be compatible with the batch environment to be used for batch production. Significant resources in terms of cpu, memory and disks are required for efficient working.
- System and Utility Software Development: The environment required is very similar to the above, but the emphasis is different. The flexibility of the system and the availability of a variety of compilers and a wide range of other tools, are very important. The quality of the Fortran system and floating point performance may be secondary considerations.
- Graphics: cpu power is important, as are the quality and degree of integration of the display. Examples are HTV, event scanning etc.
- Data Bases: The quality and ease of use of the package(s) used are of prime importance, since many users will do no more than use the database system and many will have little experience of using computer systems. Good database facilities should be made available on any interactive system.
- Computer Aided Engineering: The quality and suitability of the packages and the quality of the display are overriding considerations.
- Microprocessor Development: An easy to use software development environment is required, for ease of installation of imported programs, and also to facilitate testing of high level language programs on the host. In addition a rich set of source management and file manipulation tools is required, together with a good connection to the target microprocessor.

- Documentation Production and Mail: Such facilities must be available on all interactive systems. Interconnection of mail facilities on different systems will be required.
- Office Automation: This is a very specialised field, but one which concerns all staff involved in planning, coordinating, or managing parts of the programme, not only secretaries or administrators. It overlaps the Documentation and Mail functions in many cases. None of the systems dealt with in this report is adequate. Fortunately it is potentially an enormous market and we can expect acceptable solutions to appear in the not too distant future, as hardware costs fall. Whether or not such systems will be integrated with other interactive facilities is not at all clear.
- Administrative Applications: Closely associated with databases and office automation.
- Special Packages: e.g. Sigma, APL.

6.4. Alternatives for the provision of interactive services

As we shall see in this section, there are three main alternatives for the hardware on which to provide these services, and several variants in each case for the software. The 3 hardware choices are

- The large central computers
- Super-minis, meaning in particular VAX
- Personal Workstations on local area networks

The systems available include TSO, INTERCOM, VM/CMS, NOS, VAX VMS, UNIX, and UNIX-like derivatives. Given the requirement exposed in the last chapter, that we should seek fewer systems, both for the users and the support groups, WG3 tried to choose just one of them.

6.4.1. Central "Public" Time-Sharing

The question of when to install a full central public time-sharing system at CERN has up to now just not been answered. No one has ever been fully convinced that we had the computer resources or the effort to do it, and the compelling reasons valid elsewhere (e.g. the need to convert from the unsupported SVS system at SLAC, or to escape the home-grown ELECTRIC system under MVT at Rutherford) have just not been valid here. Furthermore, it is generally true that these other labs have had more computer resources per physicist than CERN, and therefore could escape some of the known efficiency problems of central interactive services by buying more capacity. Apart from efficiency there are other problems which are well known here and elsewhere. How to share out the resources? How to control greedy users who kill response for the others? How to maintain response and adequate numbers of access lines as the demand grows? If one looks around one often finds that resource control can be expensive in highly qualified manpower, in that elegant schemes have had to be devised and imposed (e.g. Cambridge University, ref. 5, IBM Yorktown Heights Research Labs. ref 6) in order to encourage social behaviour. Other centres, not having the manpower, were forced to take simpler measures such as limiting the time or memory size of the interactive user, so that the facilities available are in principle fully interactive but in practice are at the WYLBUR level.

It has often been considered as an approximate rule of thumb that many fully interactive systems use about three times as many cpu seconds per user per logged-on real minute as the semi-interactive WYLBUR system, which supports the belief that the latter is cost-effective. This is less true today than when we installed WYLBUR. Firstly the WYLBUR EXEC files being run today, which account for the great majority of the commands issued (≈ 95%), are of a much higher level of complication than when we installed WYLBUR. Secondly, the true comparison can only be made by including the majority of express and short jobs in the WYLBUR cost, since these jobs are an essential component of the service, and would be performed interactively if the system allowed them. One tends to find also that interactive systems which have been added to batch systems e.g. TSO, are less effective, and more costly in resources, than systems written to be interactive right from the start e.g. VM/CMS.

In terms of end-user service, one must admit that many central time-sharing installations have not been too convincing, especially if one takes sites where insufficient resources have been installed for the problem. Our "WYLBUR plus excellent debug turnround" service has compared very favourably in the end-user's mind with other high energy physics sites. The situation is changing. Users who have used interactive systems on other machines, are beginning to see the weak points of WYLBUR, and to demand full interaction on the central computers. Furthermore, the major cause for disappointment in earlier central systems, namely optimism on the number of people that could be supported on the hardware that was available, is tending to disappear as the hardware becomes more powerful, and people's expectations become more realistic.

One other problem that is growing in importance concerns the needs for privacy and security. These needs are obvious in managerial, personnel and medical applications. They are clear where we have contracted legally to protect licensed software. Whilst the majority of the physics is open to everyone, there are examples where physicists want to protect their latest papers or programs from other users. Our WYLBUR/MVS system could only satisfy these needs via the installation of a major new package which has been considered as undesirable for reasons of manpower and efficiency.

Lastly, in this brief summary, one should include some specific advantages of central time-sharing systems. One is that the resources taken during prime shift for interaction are available for batch processing use outside the prime shift, and hence for the greater part of the time. A further advantage is that users can obtain access to a very powerful cpu. This cpu, which is substantially more powerful that that of a VAX or personal workstation, is available to a single user's task if needed, provided that other tasks do not need it at that precise moment. This is of course nearly always the case on a system which is not fully loaded, or outside the prime shift.

6.4.2. Interactive Systems on Super-Minis

An alternative method to provide interactive computing in a general way is already illustrated on-site on special systems for ERASME, Merlin, LEP Database, and CAD/CAM machines, as well as several private physics group machines. In these cases a VAX computer, in the super-mini 32-bit class, is used to provide a **private** service for a modest number of users. This is a solution that clearly works well in small closed communities, although there are wide differences of opinion about how many people can work simultaneously. Suggested numbers range from 5 people doing serious interaction, to 30 people who are just editing. Although the answer depends on the size of the cpu and memory, and also upon the work being performed, WG3 has taken 10 to 15 as realistic numbers.

The question is then, if we envisage trying to run a **general** service on many VAXs situated in the computer centre, how much work is involved to translate the private success into a general one? Unfortunately, there are a number of problems. Firstly, a full service for say 150 users would involve 10 to 15 machines. Secondly, one has a number of problems that do not appear in a private installation shared amongst colleagues, such as response time control, in that you no longer know the users with whom you are sharing, and cannot go next door to complain if they are greedy. A further problem in this class is how to get all the files for a given anonymous user on to the machine to which he gets connected on a given day. In the private solution they are already there. The next problem is how to allocate 15 users out of the 150 to each machine in order to share the load reasonably across the different machines. Depending on how many greedy users are on your machine, you may get good or bad response, and this is one of the reasons for the wide spread of estimates of how many users one can support on a VAX. Next, there are a number of system utilities, necessary for general usage, that are unnecessary, (and indeed not available), in the private case, e.g. file base clean-up, control, general accounting etc. A final problem, applicable to both the general and private VAX solution, is that one must doubt whether people would make the effort to use the cpu cycles outside the prime shift. One must be careful however, in that the costing of the VAX takes this into account to some extent, so one is not necessarily obliged to use all these cycles.

A further divergence at this point is the choice of operating system for the VAX, in that whilst the Digital Equipment VMS system, which has very strong FORTRAN facilities, has been extremely popular amongst the physicists, the Bell Labs UNIX system, which is transportable to very many machines, and is thus emerging as a major World influence on everything from microprocessors to IBM (the last is not well done at present), is a major candidate on a VAX, and has strong support from professional programmers. UNIX has other advantages too, not expanded here, and a major disadvantage to date, which is its poor FORTRAN. Nonetheless it is improbable that CERN could escape having both VMS and UNIX.

6.4.3. Personal Workstations

A major new element has entered the game, namely personal workstation computers that can be connected on Local Area Networks (LANs), and hence, via a wide area network, to other LANs and to the computer centre. Specifically, such a personal workstation is a machine which today has the power of the order of a VAX 750 (half a VAX 780), has a megabyte of real memory (APOLLO can have up to 3.5 MBytes), has a local (Winchester) disk of many sizes possible between 10 and 150 MBytes, has virtual memory, has a high resolution bit-map graphics screen which gives superb graphics[7] of a quality unattainable with a conventional terminal and a communications line, and an ability to be networked together in such a way that all files on at least the LAN are obtainable in a transparent fashion. It is perhaps not profitable to continue a major sales talk in this document, it should suffice to say that one has only to look at a demonstration on the ICL-Three Rivers PERQ, or the APOLLO DOMAIN network, to see the tremendous possibilities of these devices.

However it is fair to say that, whilst these microprocessor devices point the way to the future, and the developments in microprocessors never cease to astound us, the current machines are really only the forerunners. Specifically, the PERQ has excellent graphic animation, coming from the fact that one is looking on the screen directly at the main memory, and that the device has "rasterop" machine instructions for working directly on that screen. The PERQ has also a puck or a pen for pointing at items on the screen (such as pop-up menus), or for drawing, which leads immediately to applications in computer aided design or drawing, and many other graphic applications. However, the PERQ at present has poor FORTRAN, having been designed as a PASCAL machine, a weak operating system, and poor networking. The APOLLO is rather the complement of the PERQ, having a robust operating system, excellent networking, good FORTRAN with an interactive debugger. The early versions of APOLLO had poorer graphic animation than the PERQ, having no rasterops, but this is rectified on the colour screen version and on the new DN300.

---

[7] The subject of graphics has not been discussed explicitly by WG3, in view of its coverage in other CERN discussion groups. Nonetheless WG3 wishes to emphasise the importance of graphics in the next decade, both for specific physics applications such as a Merlin successor, and for the human interface in general.

APOLLO has a touch panel built into the key-board in addition to a puck or mouse. A similar workstation is available from Xerox,[8] known as the STAR or the 8100, which is sold as a complete office system and which is not really suitable for programming. It illustrates, however, a further possibility of these devices in that it displays the commands that are possible at any point as little pictures, or "icons", at which one may point to select the command, and this improves the human interface tremendously in that one no longer has to remember obscure command abbreviations. The point is that none of these devices today on its own contains all that one would want, but taking the best bits of each one has a device of tremendous power and possibilities. In general one could say that the hardware is in advance of the software.

The price today is not right either, since the cost to CERN is in the range of 50/100 KSF per station (see also chapter 9). However the speed of entry of new devices on the market is tremendous, and one has good reasons to believe that the big manufacturers like DEC and probably IBM[9] will enter the competition to follow the recent announcements of HP. The greater the competition, the lower the prices, as one has seen in the new APOLLO stations. The recent announcements by APPLE of an entry-level PWS for less than 10,000 dollars, (and their promise of a stripped-down version in the 2,000 to 5,000 dollar range), as explained in the PWS section of the costs chapter, could have a very significant effect on prices.

WG3 believes that whilst the time is not ripe to invest heavily in personal workstations today, and indeed may not be so for a couple of years, it is extremely important to start **now** to get experience of these new devices so that we can start to understand how to use them, and to start to learn how to solve the interconnection problems that they pose. This means that money must be found to purchase say 3 stations now, and a few more as appropriate in the next year or so. We are talking of the order of 200,000 SF for each year, but it could be less with discounts. APOLLO and PERQ are the obvious candidates

[8] The Xerox Palo Alto Research Centre, Xerox Parc, was the place where such Workstations and LANs (Ethernet) were first conceived.

[9] IBM has announced a laboratory instrumentation device that has most of the components necessary for a personal workstation. It is surely now a question of marketing and presentation.

today, but we know of many others that are coming available, and the choice next year may be different. We should not worry too much that we are committing ourselves forever to a specific brand either. The money spent should be regarded as an investment in the learning curve, and we should only worry about a committing choice when we start to buy in bigger numbers.

A number of applications have already come forward as good candidates to be tried on these workstations, namely, computer aided design, for which many packages are coming on the market geared to workstations, accelerator control, experiment control, etc. However the major area to be investigated is how such machines can be used for program development and debugging, which includes large FORTRAN production codes used for physics analysis. How easily can the development be done on the workstation, and the production run in the computer centre? We already know that R.Brun has succeeded in installing all of his major packages such as HBOOK, HTV, GEANT etc. on the APOLLO, and of his conviction that the workstation speeded up his work, compared with the tools with which he was familiar, by a considerable margin. We have recently succeeded in installing essentially all of the CERN Program Library on APOLLO. Can we also show similar benefits for those programmers in the physics groups that use the computers less frequently? Are these machines easy to learn to use, and hence do they remove some of the users' problems of changing operating systems that we have discussed above?

6.5. How to choose?

Having summarised some of the options, we come to some difficult choices. It would be extremely difficult with the present staffing levels to do all of them. Nor is it necessarily desirable from the user point of view to throw more systems at him. We have seen that WYLBUR has helped us to get half a generation out of phase with some other centres, and that whilst the personal workstation concept is the first new element really to excite people in this evolution, it is just not ready today for us to buy hundreds of them.

Nonetheless there is a demand for interactive computing, strong within the specialist programming communities, and growing amongst the community of what has been termed the "DST artists", e.g. HTV users. The non-physics divisions have a wide range of programs for engineering design, administration, data bases, small scale stock control, medical records all of which are suitable candidates for interactive work, and which are able today, by virtue of their modest size and demands, to hide on the front-end CDC computers.

So the question is, do we have to provide a general time-sharing service straight away on central machines, or super-minis, or should we delay further and invest our effort into being ready to enter the personal workstation field in a couple of years' time? If the answer is the former, how should we do it; extend the WYLBUR/TSO service, invest in VM/CMS, buy larger CDC front-ends and do NOS, buy several VAX's and run both VMS and UNIX?

WG3 was unable to answer this question in the Interim Report, but has concentrated on it during the past months. The enthusiasm for the PWS solution led us for a long time in the direction of choosing this as the only solution for the interaction, together with good connections to the computer centre for batch production and file base support. Our detailed considerations have shown us that this is naive, or rather that the timing is wrong.

Costing enters into the choice as well. We have attempted to show in WG3/cba/82/32 the costs of the various interactive services. The central timesharing is probably the least expensive per user, but gives you least in terms of human facilities, i.e. limited interaction and slow, simple graphics. The super-minis cost more per user, can give good interaction and better graphics but they are still slow. The personal workstations may cost more again but give superb interaction and fast, high quality graphics, of the sort that has proved invaluable to UA1, UA2 on Merlin. You get what you pay for, and WG3 feels that there is a quantum jump in facilities between the last two steps. The problem again is that in the intermediate years, before one can have large numbers of workstations, you need to provide for those who do not have one.

The interested reader is invited to refer to the paper of Robertson, (WG3/lmr/82/45), for a further discussion, in which three possible scenarios are outlined:

- Assume that the future lies in personal workstations
- Assume that the future lies with IBM, and that means VM/CMS
- Assume that we cannot avoid a mixture of systems

We shall see in the next section that WG3 has concluded that the only realistic scenario is the last of these three.

6.6. So what should we do now?

We have emphasised that PWSs on LANs promise to provide the best way to achieve interactive computing and to improve the human interface in a way that will raise the efficiency of the users. We have indicated that they are too expensive today, and this is expanded in the chapter on costs where we make some guesses that show that there could be around 150 PWS at CERN by 1987. The correct action at this point is to form a Personal Workstation Project, which must study how best to use these machines, and how to interconnect and integrate them. This must be coupled with the purchase of small numbers of stations now, and some more during 1984. Major applications should start now with the PWS approach, and indeed we are seeing benefits already with the small number of programs that are being developed on our first APOLLO.

We have indicated the success of the VMS service within small groups, especially physics groups, and we note that this success has been achieved with very little support in most cases. Neither has there been any active "sales campaign". The success has come though the quality of the system, and its obvious suitability for this role. Many groups are buying their own VAXs, and we agree that this is a perfectly valid solution to their problems in many cases. Often it provides compatibility for the physicist between his human interface computer at CERN, his data-acquisition computer, and his computer in his home university. We recommend that CERN should support them up to a certain level. Program libraries and interfaces to the CERN networks should be provided, and their needs for file-base support must be considered. The system support, on the other hand, should be provided as it is now, from within the group who owns the VAX. Whilst we recognise the desirability of providing central systems and user support services for VAX systems in general service, WG3 has been forced to conclude that the overall manpower difficulties oblige us to leave this particular area to the client users themselves, given that this seems to work now.

We do not consider it sensible to buy, and install in the computer centre, many (meaning of the order of 10 to 15) VAXs to be used as a way of providing the general interactive service, because of the other problems that this brings, which are outlined in the previous section. The complication of the general service would remove one of the advantages of the private solution, namely simplicity.

Whilst the above two solutions fit well into our WG3 model of central batch and distributed interaction, there remains an important question that must be clarified in detail. It concerns the extent to which we need, in any case, to provide for terminal access to the central computers in the future. There are a number of points:

1) Users of distributed computers. Even if all interactive job preparation could be performed on distributed computers, users would still need some form of interactive access to the central computers for aspects of e.g. job submission/enquiry/retrieval, file base enquiry/maintenance and most probably for central mail facilities.

2) Less privileged users. The WG3 concept of fully distributed interaction will not be achieved for all users even in the medium term, given that 1400 users/week log in to WYLBUR and 1000 of them submit jobs to the IBM alone each week. We are not talking so much about the professional program developers who work all day at a terminal, but rather about a wide range of people throughout the laboratory, many of whom are less privileged or occasional users with small tasks. These people will be amongst the last to have access to a special system or a PWS, and we have to cater for their needs from the central facilities for a long time to come.

3) WYLBUR functions. It has become increasingly clear in the last year that WYLBUR plus batch is coming to the end of the period when it is acceptable as an alternative to an interactive system for the majority of users. More and more users are beginning to use interactive systems on other machines, and to complain that we do not have a real central interactive service. If one looks ahead just a few years, it seems clear that the WYLBUR solution will not be acceptable in 1985.

4) WYLBUR is coming to its limits. An additional problem that we face is that WYLBUR, as it runs now, is rapidly coming to its limits. As an illustration, the jump from 170 to 180 simultaneous users that has occurred since December 1982 has forced us to take 2 actions. Firstly, we have had to dedicate the whole of the Siemens 7880 to WYLBUR and associated tasks during the prime shift, and to stop all batch work including X and S jobs in that machine. In this arrangement the average cpu usage of WYLBUR/MILTEN, MVS/JES2 and network tasks is 80%. This leaves very little to cope with peak demand. Secondly, we have been forced to detune EXEC files in order to restore acceptable editing response. One cannot go too far down this road since EXEC files are a necessity for all users, being their present way of getting a primitive interactive service for some applications. There are other limits too, e.g. WYLBUR has an unavoidable built-in bottleneck, being a single system, in its access to the disks, as all users know from the infamous 'QUEUED' messages. The effort required to maintain WYLBUR in its present state is non-negligible and a growing problem. It is a non-standard product and its interfaces with important changing parts of the main operating system (e.g.JES2) require regular detailed updating. In summary, WYLBUR cannot be expanded more than trivially without a major development, and even then it is not clear that we should allocate more resources to WYLBUR, given that the other component of the WYLBUR service, namely X and S jobs would have to suffer the consequences.

Thus we have come to a crossroad where we must chose whether to invest in a major WYLBUR development or a new system. WYLBUR has served us extremely well but it is non-standard, semi-interactive, and limited in its design. The alternative systems, NOS, INTERCOM, TSO and VM/CMS are compared in the paper of Robertson (WG3/lmr/82/45). In simple terms, TSO and INTERCOM are ruled out on the grounds of heavy resource usage, limited facilities, and the fact that they are both old add-ons to batch systems. (See also the TSO versus VM/CMS discussion, WG3/jng,cj/82/43). The NOS solution compared favourably with VM/CMS some years ago, but we consider that it has been surpassed by the latter nowadays.

VM/CMS is hardly new, but at least it was conceived explicitly for interaction and is reasonably efficient. It is a standard IBM product and one that is very much alive and being developed as the latest release (release 3) shows. It is certainly much more complete today than when SLAC started to install it. Hence, we have chosen VM/CMS as the best way to provide a general service for those who do not have access to better forms of interaction, and as a general access system to the computer centre. We shall consider in the next section the consequences of that choice.

6.7. The particular case of VM/CMS

There are 4 levels in the possible installation of VM/CMS.

Level 1. The first and smallest effort is as a systems development tool, whereby some of the machine time at present given to systems development is available to production through the special ability of the VM system to run production and development simultaneously and safely.

Level 2. The second level is to run, as well as the former, some specific application packages for specific groups of users, e.g. PROFS (an office automation system), an electronic CAD system, or some database management system. There is a wide range of such packages available, and many of them run only under VM.

Whilst it is relatively easy to install and to support these first two layers, in that we are talking about controlled subsets of users and fixed application systems, this is not what WG3 means by a general interactive system. The real choice comes between the next two levels.

Level 3. The first of these, involving a sizeable step in effort is to run VM/CMS as a general interactive system and leave the present MVS system to run the batch (as at Rutherford). However this is somewhat inelegant for the users because of differences between the two systems.

Level 4. The final level, which involves the greatest work, is to convert entirely to VM/CMS for both interactive and batch as at SLAC. This is more elegant but requires a very serious development effort.

WG3 sees the need for an evaluation phase, lasting perhaps 6 months, in order to investigate a number of important questions. Firstly, we need to find out how complete a system, from the point of view of the users, we can put together from the latest IBM versions and the work developed at SLAC, Rutherford and other VM sites. We also need to find out how well, and how efficiently we can make VM/CMS and MVS/WYLBUR co-exist. We need to install VM/CMS such that, during the evaluation phase, the support staff have the necessary opportunity to work with the system and gain experience. We should be careful about level 2, in that packages once installed are very hard to remove. The third level should only be tackled after a successful outcome of the evaluation phase. Level 4 is more ambitious, and can only be tackled in the long term, however desirable and more elegant than level 3.

One specific reason for the evaluation phase merits some further explanation. VM/CMS, as it is supplied by IBM, is not a particularly user-friendly system. However it has excellent potential in its design for being greatly improved in this aspect, in that each user can build his own commands and facilities; in short his own personalised environment. It would improve greatly the acceptability of the system at the beginning if a good standard environment and set of tools were provided centrally. Failure to do this has led to negative user response to VM/CMS in other sites. Other justifications for doing this work include catering for the infrequent user who will never have time to develop his own working environment, and the desirability of not having each user working in a system that looks different to each other user. This is precisely the same approach as we adopted in installing the successful WYLBUR/batch service in the past. Nonetheless, no group would be prevented from creating its own environment if it so chose, using the tools that VM/CMS provides. A major objective of the evaluation phase is to improve the user interface of VM/CMS for CERN users, or as a minimum to understand how much work is involved.

A consequence of entering into VM is that WYLBUR would necessarily be constrained to a minimum development, and ideas to develop a full-screen WYLBUR for CERN must be dropped. This does not mean, however, that the service given to WYLBUR users should decrease. In fact it is important that the service be kept at at least the present level for some time, and this requirement implies extra costs. It is probable that we shall have to find a way to make a small increase in the number of WYLBUR lines available, because of the continuous growth in demand, before we have a VM/CMS service sufficiently mature that it can support a significant number of users.

We have chosen to recommend an initial evaluation of VM/CMS because the best way to find out the latest state of the product is to work with it, rather than to listen to a certain amount of existing historical prejudice. We need to purchase some hardware in order to do that without interference to the present WYLBUR. **We should emphasise that this hardware will be needed in any case, whichever route we finally follow, in order to cope with the inevitable short/medium term expansion of facilities on the central computers.** We shall return to this point in the chapter on costs and manpower.

7. TRAINING

This topic, although short, merits a chapter to itself because of its importance. Training courses in computing have been felt largely unnecessary in the past, except in specific areas, and it is only recently that it has become clear that we need to set up training for three groups of users for whom there are presently no courses.

Historically, the major part of the user community consisted of physicists, or rather the physicist computing specialists in each group. Currently, with the number of computer users having doubled in three years, we have not only the physicist specialists, but a much greater fraction of each physics group using the computers. Furthermore, the number of non-physicist users has greatly increased, meaning engineers, administrative staff, secretaries, etc. The only CERN training courses offered today consist of the Technical Training programme, in which courses on advanced WYLBUR, INFOL and SCRIPT are given at the level of junior programmers or data-aides, and a few ad hoc courses given by the programming enquiry office on request. A purchased self-study course on IBM job control language, using audio and video tapes, has had almost no interest, despite the obvious need for something on this subject. The computer operations group has a very small and necessary internal training section which endeavours to keep the operators up to date, and system programmers have attended (of necessity) the courses offered by the manufacturers.

At the Software Workshop in October, (refs. 2 and 3), a number of points became quite clear. Firstly, the support required by engineers, and administrative staff was illustrated by the talk of Turnill, (who used to be at CERN, and is now at Britoil). These users require a different sort of packaging and training from that accepted in the past by physicists. They demand fully finished application packages with a good human interface and a flexible command language with which they can achieve their work, without having to program in FORTRAN. This is quite different from the style of the physicists, for whom FORTRAN is essentially the only tool they know, and who request packages which they can incorporate in their own FORTRAN programs. The engineers demand training courses in how to use the packages, the physicists have traditionally requested documentation rather than formal courses, and expert consultation when things go wrong.

The feeling has always been that the physicists do not want to waste their time going to formal courses. They will rely on natural intelligence and in digging the appropriate part out of the manual. There are signs that this is changing. The two lectures by M.Metcalf on FORTRAN 77, repeated this year, had a very good attendance from physicists, and are indicative that we should revise our policy that, because we think the physicists don't want courses, we don't very often offer any. In conclusion of this point, we recommend that training courses are offered both for physicists and for non-physicists. That being said, there is a fundamental problem that we are not staffed, either in numbers, or in training, or in style to offer such courses. We enlarge upon this point in the section on manpower in chapter 10.

The second point that became clear at the Software Workshop, concerns the computer professionals themselves (including the specialist physicists). The world of computer science, regarded somewhat scathingly in the past, is coming of age. New concepts and techniques are emerging that could well be appropriate to our work, and help us to stop reinventing our own wheels. It was appropriate in the past for high energy physics to develop their own tools, and indeed we have been very successful in that many of our tools stand close comparison with those available in other scientific fields that make heavy use of computers. At the Software Workshop we saw that we need to open our eyes, and to investigate the offerings of computer science. Unfortunately, we are so inadequately educated in this area that we are not able to make even proper evaluations without some training. This point is expanded in some detail in the ECFA report of the workshop (ref. 3). WG3 fully supports the recommendation on training contained in this report, namely:

- "A person or small group is put in charge of introducing one or several courses in CERN's Academic Training Programme with titles and speakers chosen to cover topics of practical application of computer science concepts. Such courses are strongly recommended to all experimental physicists, and pushed even more with professional and semi-professional software experts. The programme is maintained or enlarged according to its success, although attendance alone should not be the criterion: Modern industry enforces technical training on all their higher level staff, thus ensuring that a formidable investment, the personnel, is not allowed to degrade".

We also support the corresponding recommendations of the CERN Joint Training Board (ref. 7) on regular professional training of the kind indicated above, together with its implication that this will require accepted diversion of manpower for a week or so each year.

8. OTHER TOPICS

8.1. Electronic mail, voice mail, and documentation systems

8.1.1. Electronic mail

A paper on electronic mail, (WG3/lmr/82/35), was presented to WG3 by L.Robertson. It shows that even the simple WYLBUR Messages facility is strongly used, as has been the older ELECTRIC Mail system in the UK. WYLBUR mail is used by more than 600 people each week, and by more than 900 in a month. It is an essential communications tool for many people including, of course, large physics collaborations. This has been clearly supported in all the discussions on external communication connections to CERN, for example in the ECFA sub-group 5, whose survey of European laboratories shows an unexpectedly large number of users of various systems, and a demand for a coherent general service.

A pilot project on WYLBUR, called EMF, giving extended facilities for distribution lists, replying and forwarding, deleting and setting a reminder date, and filing with keys, additional comments etc. has been a success. We have had doubts whether it should be released in its present form, because its implementation in EXEC files is certainly not the ideal way to do it. Should it be implemented in assembler within WYLBUR, implemented on a special MAIL machine like a VAX, (would people login to a special mail machine?), or put under VM/CMS which would give a sounder implementation, as well as some possibilities for privacy which are necessary for some aspects of electronic mail?

Many more details are given in the above paper, covering important questions such as how to link-up the many mail systems that will (and already do) exist. They will not be expanded further in this report, but suffice it to say that WG3 believes good electronic mail facilities to be a very important tool and they must be provided.

Our conclusion is that, even if we start a VM/CMS service next year, the majority of users will be working with WYLBUR for some years, and that the EMF pilot project should be tidied up in the light of experience of the trial phase, and released for general usage. A version with the same external specifications should be provided under VM/CMS using the same message format as far as possible, and facilities for passing these messages to other machines, on and off CERN site, should be gradually developed. We are aware that this is far from the ideal solution for many reasons. Unfortunately we see no other practical alternative given the present demand, not least that coming from the LEP collaborations.

8.1.2. Voice mail

Voice Mail is a relatively new logical extension to electronic mail. Here the telephone is used rather than the terminal, and the voice messages are recorded on a computer disk. One needs either touch button telephones, or a cheap (30$?) attachment that clips on the handset, and one obtains many of the facilities listed above, except that one gets a spoken rather than typed message. These systems are available on the market today. Given the lack of success of contacting people via the CERN telephone system, one can only support the tremendous time-saving facilities of such a voice mail system.

Voice Mail and Electronic Mail are complementary in the sense that, whilst the former will reach a much greater number of people, (we don't all have terminals and use the computers regularly), and is ideal for short messages, the latter provides for exchange of electronically stored material for programs or documents, has much greater storage capacity, and facilities for indexing and searching.

The costs of voice mail seem not to be high, although a major unanswered question is whether the CERN telephone exchange itself is too old to be used with such a system. If it can be used, then the cost for a 1000 subscriber system from any one of the 3 manufacturers advertising such systems that we have seen, (there may well be more), is of the order of 150 KSF, plus the clip-on handset attachments. WG3 has not had time to make a detailed study, indeed the topic is very new, but we recommend that a proper study be made. We have certainly noted a great deal of support and interest for the idea, especially from visitors trying to make contacts with people at CERN. The savings in time could be very significant.

Finally, we should note that it is becoming apparent that CERN must replace its telephone exchange on the timescale of LEP, and a digital PBX can offer voice mail as a standard feature.

8.1.3. Documentation systems

A final topic in this section, which we have not explored in detail, concerns electronic document filing systems. Here video disks have tremendous potential, and we are waiting for developments. All we can say at this stage is that many people, including the participants in these working groups, are asking for a fully supported document filing scheme at CERN, and that WYLBUR as it stands poses many problems for a correct implementation. In particular, many large physics collaborations including the LEP experiments, are demanding such schemes with some urgency.

The problem is not only that WYLBUR is inappropriate for the full task, again because of EXEC file implementations, and lack of privacy, but that there are at least two other questions to be answered. Firstly, there is a technical problem in that computer systems in general have not solved the problem of how to store all the components of a document, namely text, equations, graphics and half-tone photographs. It seems that, apart from the many computer science departments in the world that are working on this problem, the major manufacturers such as IBM and Phillips are investing heavily in finding solutions. IBM has recently announced a document scanner as part of its office equipment line. Secondly, there is a real problem in deciding what people would like to see as the specifications of such a system. The EMF scheme mentioned above has one form of text file base built in, and many other implementations exist, even on WYLBUR, such as that used by UA1, or the scheme of WG2. None of them fulfils all of the needs, and we are at a loss to know what to recommend, given that different people are prepared to make different compromises. At this stage, this seems to be a very personal matter, but a consensus must be found soon.

8.2. Languages, libraries, PATCHY+, and database systems

All of the topics in this section have been discussed in sub-groups who have produced reports for WG3. Two of the sub-groups already existed within the framework of the ECFA computing studies, and since they contained many of the relevant people for the topics in hand, WG3 chose not to duplicate their work.

8.2.1. Languages

M.Metcalf formed a small group looking at the impact of languages in the next decade. They have produced a report (WG3/mbm/82/39), which has been agreed by WG3. They make one simple clear statement. As far as physics analysis programs for LEP are concerned there is no alternative to FORTRAN 77. This was confirmed by the interesting discussion on this point during the Software Workshop in October (see CERN 82-12). The real choice facing the physicist community will come only later when ADA, (from the USA Dept. of Defense), and FORTRAN 8X become available. WG3 recommends that those people at CERN, who are already following closely the developments of these two languages on a voluntary basis, be supported in their work.

8.2.2. Libraries

H.Renshall, both in his role as CERN program librarian, and through his membership of the ECFA computing sub-group 7, has made a report to WG3 (WG3/hrr/82/55) on the subject of public program libraries in the LEP era. The interested reader is invited to read the above paper for full details.

The program library at CERN has just been completely restructured into what might be described as a 5-layered shell. This was done in response to the need for physicists to be able to get the programs they want from one organised structure, rather than the several different and independent program libraries that had grown up at CERN. It was also done in association with the change from FORTRAN 4 to FORTRAN 77. It is worth pointing out that, whilst the inner 4 layers of the library are fully supported, the outer section, consisting of physics algorithms, or codes that might be of general use, perhaps after specific modifications to the source, e.g. the LUND Monte-Carlo, has really no manpower to support it, and is very much a spare time or goodwill exercise. However, overall, the library at CERN is now in a rather good state, and versions exist for CDC, IBM, and VAX. Preliminary versions have been made for APOLLO and NORD-500.

One of the real problems that exists in trying to make this library available outside CERN, in a consistent way, and at the latest level, is that few of the outside institutes cooperate at the level of their computer user support groups. We are still working mainly in the traditional way, where one (or several) physicists from a collaboration collects a tape from the program library and installs what he needs on his home computer, independently of any other physicists in the same institution. Several different versions of the library end up on the same computer with no coordination between different experiments. Furthermore the library staff have to make more copies than are necessary. We have taken steps to remedy this in two of the larger labs, namely Rutherford and DESY, by going there and setting up a collaboration with the computer centre management, so that one of their staff is named to be responsible for the CERN library. WG3 fully supports the extension of this approach, but emphasises that it can only come about if pressure is brought to bear upon the computer centre staff, by the client physicists themselves or by ECFA.

A final point is that, beyond the above step of making the CERN library available in all visitor's institutes, there is a natural request to make a common (European) high energy physics library. However desirable, this is unfortunately out of the question with the present manpower available, at least at CERN.

8.2.3. PATCHY replacement

PATCHY has become invaluable as a **machine independent** source code maintenance and transport facility in this field where individual maufacturers' systems are useless for large collaborations using multiple machines. It has however the drawback that the source maintenance facilities interfere with the editing and debugging facilities that exist on individual systems, as was pointed out many years ago by M.Metcalf in a paper to the now defunct Computer Users Advisory Committee (CUAC, ref. 4). Hence an extension to PATCHY is being discussed by a small group lead by H.Grote. A report (WG3/hg/82/34) has been written, and discussed by WG3. We find the general line of approach to be interesting, but that we were far from able to agree on the details. Furthermore there are many interested parties not in WG3. We recommend further studies of Grote's proposal with a wider group of people during next year. However, we feel that they should find out more about possible replacement systems outside CERN, for example one of the components of the TOOLPACK project, which is being created by a cooperation between several organisation, (Argonne Nat. Lab., Bell Labs., Jet Propulsion Lab., IMSL, NAG, Purdue Univ., UCSB, Univ. of Colorado), and being funded by the National Science Foundation and the U.S. Dept. of Energy. The documentation available to us on this is insufficient to see whether this is a real competitor for a PATCHY successor, but if it is, it is clearly better for high energy physics to adopt what could become a much more widely installed system, just through the weight of the above supporters. A further example, named CCC, appears to come from a company with the unlikely name of SOFTOOL.

8.2.4. Databases and Bookkeeping for HEP

WG3 chose not to study this point in detail because of the existence of ECFA Subgroup 11 on "Databases and Bookkeeping for HEP Experiments". This group is chaired by John Hart of the Rutherford Appleton Laboratory to whom interested readers should apply for the working papers. The aims of the group in summary were: to provide a guide to the packages in present use, to find out the future requirements, and to make recommendations how these can best be met.

John Hart presented an interim report to WG3 in January 1983, but the recommendations of the subgroup were further discussed and presented to ECFA at the Amsterdam meeting in March. It is hoped to publish the report in May.

One of their main recommendations is that, on the timescale on which software is written for LEP (1985 onwards?), relational DBMSs should be used for all but the simplest database applications. Only relational systems offer the flexibility and ease of use required for HEP. The development and maintenance of HEP packages can be costly in terms of manpower. The query language SQL is becoming a de facto standard for interactive and program access to relational DBMSs. By adopting this as a standard interface, it would be possible for the groups within a collaboration to use the different DBMSs available on the computers at their institutes. The use of SQL would generally remove the need for special purpose listing and editing programs.

There is a problem of price whenever one uses commercial systems, but this may not be as serious as originally feared since less expensive products are already appearing on the market and the approach outlined above would allow DBMS systems to be shared with other scientific and administrative applications. There are no special HEP requirements for a relational DBMS which would by themselves justify the very considerable manpower effort (at least 10 manyears) needed to write one specifically for HEP use.

Subgroup 11 also point out that relational DBMSs can be used to store any tabular data, and therefore have applications other than bookkeeping within HEP. For example, a DBMS could provide a powerful catalogue facility for other data, such as bulky calibration data, which is stored separately. The existence of a Fortran interface means that it is simple for another package to communicate with the DBMS. Thus a package to store and retrieve calibration data or graphical information could readily be constructed.

In the short term (for the next 2 or 3 years), there will be a continuing need for simple HEP database packages. This should be satisfied by the upgraded version of KAPACK which is now available. At the same time, HEP groups should be encouraged to try using relational DBMSs for suitable applications to evaluate their usefulness and build up a pool of experience with them. For this reason it is important that relational DBMSs supporting SQL (eg IBM's SQL/DS, ORACLE or perhaps RAPPORT) are made generally available, preferably on the central computers, at the major HEP centres.

An additional point, put forward in WG3, is that HEP should build specific database software behind the new memory management package GEM, and thus use GEM itself for database-like requirements. The arguments in favour of this are not only that it avoids the above pricing problem, but that many data structures used in GEM are related to the data structures needed for databases. The reader is invited to see the minutes of WG3 meeting 20 for the discussion on this point. A proposal has been written as an internal DD document, and the interested reader is invited to contact R.Brun for information. This alternative approach brings us right back to the arguments familiar to the attendees of the software workshop i.e. should we go it alone or buy it from outside?

8.3. Computer Aided Design, and Micro-processor support

The Euclid system for mechanical CAD has been chosen and installed recently at CERN on a dedicated VAX. Computer aided design for electronics has been under study for about nine months in another committee at CERN. WG3 has received their status report made to the DD Joint Planning Committee (WG3/mfl/82/48). We note the difficulty in making a choice at this stage of the game, especially in finding an integrated system from one manufacturer. Since they have not produced any firm conclusions we can only support their general approach. One thing is clear: WG3 believes that without suitable CAD tools, we shall not have access to the latest electronic technologies, and we shall be unable to design some of the electronics needed for the next decade, e.g. semi-custom or custom chip design, and we will be left with outside contractors as the only solution. In addition, large dense boards such as those required by FASTBUS will require more powerful design tools than are available at CERN at present. Finally, electronic CAD systems can have a major impact on the productivity of the layout technicians for the printed circuit board work we are already doing.

Micro-processor support is yet another area where WG3 has felt not particularly competent to get involved, and indeed there was no need for their involvement, given the other committees already working in that area. We note the recommendation that central micro-processor support be provided on a VAX, using the UNIX operating system, in that it is another operating system, another console and another set of users to support from the central pool. Nevertheless, we fully understand and support their recommendation.

9. COSTS AND MANPOWER

9.1. COSTS

9.1.1. Batch processing power

WG3 recommends, as a minimum, the doubling of the batch processing power in the computer centre by 1987, in time for the start of LEP experiments. It also points out that it is possible to install this in an incremental way as needed, with the proviso that the capacity must be installed in time to avoid the very negative effects of saturation. There is a rule of thumb that has applied to computing in the last decade that, roughly, one can purchase twice the power for the same number of francs each four years. Hence, globally, a doubling by 1987, should be attainable at the current level of budgeting of the computer centre, if the rule of thumb remains applicable over the next five years.

In practice the situation is very much more complicated, due to the mixture of rented and purchased equipment presently installed, the need to install not only cpu power but major peripherals such as mass storage upgrades, disks and probably tape units. CERN's particular "shop window" position can lead to prices considerably different to catalogue prices. There is the possibility that at least one LEP experiment contributes a major computer which could be sited in the computer centre. It is thus very difficult to make a precise plan of how all this will fit into the budgetary envelope of the next 5 years.

In the whole of this chapter, all numbers are given in 1983 prices unless explicitly stated otherwise. Looking at the past three years, we have spent an average of 14.3 MSF per year on central computing, if one adds the operating budget to the capital spent on acquisitions, (excluding networking and terminals). This represents 4% of the CERN annual materials budget. The CDC upgrade which is under negotiation should add of the order of 1.5 to 2.0 units of IBM 370/168 power for 1 MSF/year less than we are paying CDC today. Having finished the acquisitions on the IBM side of the room, and renegotiated the CDC contract, we will be at a 1983 budget level for pure operation of 9 MSF, which is 5.3 MSF/year less than the yearly spending over the past 3 years. This budget level allows for no additions or improvements whatsoever, and it is clear that it will have to be augmented if the new needs of the user community and of the physics programme are to be met.

Allowing for the increase in CDC capacity, we will still need to add at least 8 IBM 370/168 units of computing, together with a number of computing peripherals, (discussed in the next section), on a time-scale to 1987. The question is therefore, how much should be installed, in how many stages, and at what times? We have seen from figure 5 that, based on a short term load extrapolation, we will be saturated in 1984. We can see that, following their recent run, UA1 and UA2 currently represent half of the IBM load. This can only increase in 1983/84 as the annual integrated luminosity of the collider increases. Fixed target physics, through improved triggering with micro-processors, is tending to increase its processing demand for a given amount of beam time. ISR physics processing will continue well after the end of 1983, when the machine stops. It is therefore highly probable that, despite the controlling influence of COCOTIME, we will be unable to wait until 1986 to install more batch production capacity. Furthermore, COCOTIME can influence only the production demand, and we have **no effective levers to control the prime shift demand,** which is also approaching the limits of the installed systems.

We cannot be sure when we will need more cpu power, but we can present a **model** of incremental addition of capacity, which should be seen as **indicative** of the scale of the investment needed. In this model, the factor 2 is achieved in 2 steps, starting in 1984 by the replacement of the existing Siemens 7880 (Fujitsu M200) by a Siemens 7890 F (Fujitsu M380, also sold as an ICL Atlas 10). The single-processor version would provide 6 IBM 168 units, whilst the sale of the 7880 would remove 2.4 units, giving a net increase of 3.6 units. A second 7890 processor in 1986 would add a further 5 units, thus giving a total by 1987 of 8.6 units which matches the model.

The single-processor, after deduction of the resale value of the 7880, could cost of the order of 5 MSF, and could be upgraded to the double-processor version for a further 5 MSF. To be specific, the M380 with with necessary minimum peripherals, described below, costs 7 MSF, and the resale of the M200, perhaps packaged with a greater discount than we have considered, is entirely a guess at 2 MSF, depending strongly on when we do it. For the above prices, we have taken catalogue numbers minus a modest discount of 15%. The true cost for CERN could well be lower via a competitive tender, or because prices tend to fall with time. The price could be higher either because we may well need more memory than has been included in the above costing, or because we have been optimistic about the 7880 resale value. Nonetheless the model could well be achieved with a smoothed annual capital expenditure starting at 2 MSF in 1984/5, and rising to 3 MSF for 1986/7.

We have included some peripherals in the price, because it is no use installing greater capacity without the minimum peripherals necessary to start exploiting it. The price of those peripherals has been taken to be 1 MSF, which could purchase for example 4 strings of 4 disks each, together with the controllers. In reality the detailed choice of peripherals would have to be made nearer the purchase date. Furthermore, other peripherals such as mass storage upgrades, laser printers, more disks and channels etc. would have to be added to balance the configuration of the whole centre, and these are mentioned below (i.e. not in the above prices).

In summary, this financial model shows that, even at today's prices, the desired minimum capacity could be installed by 1987, in two stages, at an entirely reasonable cost of 2 MSF per year in 1984/1985, and 3 MSF per year in 1986/1987. These costs are sensibly less than the 5 MSF per year that would keep us at the budget level of the last 3 years, but this is far from the whole story. We have included only the bulk capacity component, and there are many other budget lines to be added in order to make a complete service.

9.1.2. Mass Storage System

At the end of 1981 the DD Division requested an upgrade of the existing IBM Mass Store from a model A2 to a model A3. This request, which was postponed, would add 50% to the number of read/write heads, and 66% to the volume of data that can be stored. It is also necessary to add another component, which, like the read/write heads contributes to the throughput of the device. This consists of an increase in the number of staging disks associated with the mass store. The total cost for this upgrade, which is in any case smaller than the upgrade to an A4 model which may become subsequently necessary, is 800 KSF. The justification is clear to all users, on many machines around the site, not just those on the IBM, in that it becomes very difficult to work when the MSS is saturated, as it has been on several days this year. Once saturated, it is in the nature of the device that a "thrashing" effect becomes suddenly apparent, and the throughput of the device drops dramatically. Furthermore, the number of tape mounts on the IBM system has risen sharply in direct correlation with the unavailability of new MSS volumes. The end result is that, since we cannot afford to have all the users waiting for hours to get their files, upgrades **must** be made in time to avoid saturation.

The MSS provides an essential service for all users of the IBM, and for many users of mini-computers on the site. It has kept down the cost of disks on the mini-computers, which has tended to be substantially more expensive per megabyte than the disks on the central computers. Today we are at the limit, having applied all the personal intervention techniques possible to optimise the usage of the major users, and hence to delay the previously requested upgrade.

9.1.3. File Server and "Hyperchannel"

We have discussed in section 5.3 the desirability of a file server computer connected to the mass store and its justification. We have explained that WG3 did not have the time to make a serious study in this area and we have recommended that it be made. Based on what we already know, we believe there is a very high probability that it will be cost-effective and probably unavoidable to install a file server machine. We must therefore include it in our cost table. The problem is that, without the study we do not know how big that machine must be. Hence we make a guess that it would probably be a machine of the size of the upper end of the 4341 range, equipped with rather minimal peripherals. The staging disks associated with the file server are already included in the mass store cost. No major communications controller would be needed. One might well include some tape units for backup, or to off-load some of the many tape copy/tape split jobs that complicate the scheduling of our large machines (a problem which is becoming serious). The decision to purchase or to rent is also difficult, since the advantages of one or the other vary with the size of the machine and with time. We end up with a guess that we will have to spend of the order of 1.5 MSF over 3 years.

The associated problem of how to connect many machines in the computer centre at speeds higher that those achieved at present requires a further study as we have recommended. Since the IBM mass store can only be connected directly (channel connection) to a maximum of 3 (IBM-compatible) computers, it is clear that we need to tackle this problem. Apart from CDC and VAX computers in the machine room already, we may well face a situation where more than one LEP experiment installs a major computer for its own purposes in the computer centre. There are many detailed questions that we cannot expand on here and which can only come out of the study. Hence the figure to be put in the cost table is again a guess at about 200 KSF. Clearly it will depend on the cost of the interface to each individual computer. This could be quite significant depending on the choice made.

9.1.4. Other major peripherals

We have included some major peripherals in the costing of the batch processing power and of the VM/CMS service, (e.g. disks, controllers, channels and main memory), but we should spell out explicitly the relatively short term requirements for other peripherals whose costs are not in either of those sections.

Currently, we have problems finding enough tape units to maintain enough concurrent jobs in the IBM machines to exploit the cpus at night and at weekends at 100%. This is explained by a certain tendency for production jobs to use more tape streams per job than in the past, coupled with the fact that we have the same number of tape units as when we had only 2 units of IBM power, (we now have about 7). The number of tape mounts per week on the IBM system was constant at about 2000/2500, from 1977 until Autumn 1981. Given the drastic increase in the number of jobs run per week over the period, (figure 7), this has always been attributed to the success of the Mass Store. In Autumn 1981 the mass store saturated, there was no more space to give out, and we installed the 3081 computer. Tape mounts rose to around 3000/3500 per week. Under recent production pressure from the UA experiments, tape mounts have risen to 5500/6000 per week, leading to sustained periods with one mount per minute. As a start, 4 new high speed IBM tape units and a controller would cost about 260 KSF.

Communications controllers present a problem of choice on the IBM side because of changing technology and access methods. We are close to the limit of the existing equipment, (IBM 3705s), and are reluctant to make another step with these rather old devices. Figure 9 shows the regular increase in the number of users per week, which is reflected in demand for terminals and lines. The negative effects of being unable to get a line when one wants to work are well known, and it should be unnecessary to spell them out here. (They are used as a standard argument by anyone wanting to purchase private computing equipment). The alternatives to spending 300 KSF on a new communications controller (IBM 3705) are either to increase the number of connections that come in via CERNET and the WYLBUR bridge, or to increase the number of Series 1 computers, (at a cost of 85 KSF each), that allow full-screen access from Pericom terminals attached to INDEX, (which is a very cost effective way of gaining full-screen support without installing a completely new set of terminals, cables and controllers). Both must be done eventually, but it is possible that neither can be achieved on a sufficient scale to avoid acquiring a further 3705. Time and the development of the demand will tell.

Apart from these specific examples, and the case of the printers that follows, one must be very clear that a computer centre as large as that at CERN, with a strong growth curve, needs a reasonable money buffer each year to cover new peripherals. These cannot be foreseen in detail, but rather become apparent as the evolution in demand reveals bottlenecks. We can be more specific about the peripherals we need for 1983, but for the following years we must include of the order of 750 KSF/year for a general peripheral buffer. With the existing growth curve we have been regularly in the position of running out of one resource or another, and it is well known that saturation in computing comes often as a sudden and drastic effect. Until recently there has been enough money in the computer centre budget to handle most of these situation by buying, for example, a couple of disks, or more communications lines. That flexibility has been removed in the present budgeting. Experience shows us that a minimum figure on the "IBM" side of the room is 500 KSF/year. On the CDC side we have been running a stable configuration for a constant user community for several years. The proposed CDC upgrade will certainly imply some more peripherals to balance the configuration in the future, given new users and the fact that, during the tight negotiations for cpu power, the peripherals were cut to a minimal level. In summary, it would be foolish not to include a sum of the order of 750 KSF/year in this budget level. Without it the computer centre will be unable to evolve the services to match demand.

Finally some words must be said about the view, commonly expressed in some areas, that computing is a vicious circle. You buy mainframes one day, and the next day you need more peripherals. The day after you need more mainframes. Indeed it seems like that, but there is a reason which is usually not stated. There is a tremendous growth in usage of computers at CERN, (see the figures in chapter 4), and it is through this growth that more people become more productive in their work. This is hardly unreasonable in a Laboratory that works on the forefront of technology. A more positive view would emphasise that the annual cost per user, and the annual cost per cpu unit have fallen by a factor 2 in the past 3 years.

9.1.5. Printers

Printers merit special attention, not only because the computer centre spends more than 1.1 MSF/year on printer rental/maintenance and the paper budget, but also because they consume annually of the order of one man-year in servicing the remote stations, and 5 man-years of the contract labour in the computer centre, (one man/shift). This contract labour translates as an additional 300 KSF/year. Furthermore, we must budget for the replacement of our remote stations which are more than 10 years old. In detail we have the four following considerations:

1. **The Paper Budget.** It is somewhat surprising, given the doubling of the number of users over 3 years, that the paper load has not doubled. The remote station usage, limited by speed of communications and other controls, has stayed roughly constant. The overspill on the central printers has risen approximately 50% in 3 years. Microfiche outputs, substantially cheaper per page than paper, have trebled during the period and have helped to contain the paper budget. The IBM 3800 laser printer has enabled the use of standard sized paper, smaller and cheaper than its "pyjama" predecessor. Advance purchases of paper have delayed the impact of increased load. The resulting possibilities for economies, which have kept the paper budget roughly constant at 500 KSF/year, have now reached their limits, and more and more draconian constraints on users will be needed in the future to avoid increased paper cost.

2. **RIOS Replacement.** The 13 Modular One (RIOS) remote input/output stations (nowadays essentially remote printers) on the CERN site are 10 years old and expensive to maintain. Three of them have just been removed to reduce this budget. We need a RIOS replacement on a rather short time-scale. No product exists today that matches our ideas of an ideal "one for one" RIOS replacement, i.e. a printer of modest price, non-impact technology, fan-fold paper. Given the short term need, one is obliged to consider a change in strategy to try to match the problem to the existing products.

3. **Printer Technology Revolution.** Impact printers are expensive to maintain, and limited in speed, quality and characters available. They

require non-standard paper. New technologies, using photocopier techniques, with lasers or ions to draw the image, are bringing changes in all the above parameters. The IBM 3800 was sold as a solution to the bulk printing problem. It enabled us to return 5 impact printers. People quickly saw new possibilities, unforeseen by the designers. (One man at least had an idea good enough to make him a millionaire). More generally, many people saw the possibility of high quality output, direct from the computer onto paper, eliminating time lost in multiple retyping, or photocomposition, or print shops.

4. **Document Printing.** The 3800 showed us the possibilities, but has marginal quality for many applications. Better quality laser printers are coming on the market at an increasing rate to meet the world demand for document production. Their particular relevance to CERN is coupled to a number of factors:

a) CERN's major direct output is physics papers. The character set range demanded by the greek and mathematics, coupled with the European languages, exceeds all impact printers and the 3800. The demand for integration of graphics and even photographs defeats all but the very latest offerings.

b) More and more professional scientists and engineers find it productive to type their own papers directly into the computer, and to be able to to edit without major retyping. They see further advantages in direct viewing of documents at a terminal and in electronic mail and filing. There are many other people don't have enough secretarial help at CERN, and are obliged therefore to type the document themselves. All these people need local document-quality printers.

In summary, we have several printing problems to solve, at a time when the market is changing rapidly in response to need.

The present market offers:

- IBM 3800, and its direct equivalents: High speed, modest quality, fan-fold paper, all-points addressable (APA) promised for a new model, which will allow mixed text and graphics in the future. Poor turnround coming from the need to run two paper sizes. Price around 800 KSF for present model, more for APA.

- Xerox 9700: High speed, high quality, cut sheet standard size paper. Large number of fonts and graphics. Saclay use the 9700 to print 4 normal pages on one recto-verso sheet with good readability. This allows a significant reduction in the paper budget. Unfortunately the cut paper creates new operational handling problems. We are looking at solutions such as shrink-wrapping, binding, or stapling, but all involve more operator effort than fan-fold paper. Price 800 KSF.
- HP 2680-A: Medium speed, high quality, fan-fold paper. Large range of fonts. Well made. In many ways an excellent RIOS replacement, were it not for the price, which is a factor 2 too high, especially since you have to buy an HP3000 computer to go with it.
- IBM 6670: Modest/low speed. High quality. Recto-verso. Cut paper. Reasonable selection of character sets (effectively 5 times the 3800). Possibilities of APA on the horizon. Reliable and robust small printer. Price about 120 KSF.
- Miscellaneous laser printers of small photocopier size: Low speed, high quality. Cut sheet. No recto-verso. Many photocopier firms getting into the market, without necessarily having the expertise in computing/electronics. Character set/graphic possibilities depending on the quality of the interface. Small paper hoppers. Unsuitable as direct RIOS replacement. Very suitable for small local private printer, e.g. local control room, corridor, local area network printer etc. Price 40 to 70 KSF, but not quite on the market yet.

In summary, the computer centre must look seriously at Xerox 9700 as a way of reducing the paper budget, or wait for cheaper versions of 3800 style machines. The RIOS replacement question is not solved today. Detailed calculations show that it is unlikely that any significant fraction of the RIOS replacement problem can be funded from the present maintenance budget. We are obliged to budget for a sum of the order of 1 MSF for the replacement of 10/12 objects at around 80 KSF each. The last category of printers offers a possibility of change of strategy: more local private printers attached to computers on CERNET would not only improve the geography, but distribute the responsibility and the paper budget amongst the users. The latter seems to be about the only effective way of making people conscious of how much paper they are using. Finally, high quality document production is a growing demand that must be satisfied.

Given all these factors it is very difficult to predict the overall effect on budgets and manpower at this stage. However we see the need for a small study group to make some decisions in this area rather soon. The budget line in the cost table provides for the RIOS replacement and a small document printer only.

9.1.6. Costs of a central interactive service

We have seen that we need to add an interactive service to our central batch and WYLBUR services, which in the long term would replace WYLBUR with a more standard IBM product. We have recommended VM/CMS. One problem in costing that service is how to separate it from the other services offered. An interactive facility would be a significant drain on resources only during the prime shift, which represents 50 of the 168 hours in the week. It is indeed one of the advantages of central interaction that the cpu, memory and I/O resources taken for interaction during the prime shift may be used for batch production at other times. A second problem comes from the constant growth in the number of people using the central computers. The new demand has to be satisfied one way or another, either using WYLBUR+batch, or using VM/CMS. We have explained that, in contrast to a full TSO/SPF service, we believe that a VM/CMS user is not a greatly higher drain on resources than a WYLBUR user, once you count the cost of his X and S jobs run during prime shift. The SLAC usage figures support this belief.

We also believe that a VM/CMS service should be introduced in steps, both in terms of numbers of users, and in terms of functionality. Furthermore, the growth should be modulated by the success of our preferred interactive systems, namely personal workstations and private VAX/VMS services. This means that the users with heavy interactive demands should start as soon as possible to use the workstation approach where the facilities are best geared to his needs. This will help to limit the demand for the central interactive services.

We present two models for the introduction of VM/CMS which differ only in the introductory phase. The first model assumes that there will be enough capacity during 1983 on the existing central machines to introduce VM/CMS in addition to the present services, at least for evaluation purposes. The second model assumes that there will not be sufficient capacity, and starts the VM/CMS evaluation on a small 4341 machine. Both models assume that a serious user service can only start after a major capacity upgrade has been made. In the second model the 4341 is reused as the starting machine for the file server. Both models present the **extra costs** of adding the new service. The first model explains in detail the need to purchase equipment such as memory, channels, communications controllers and solid-state disks. Whilst these are associated in the text with the VM/CMS service, the greater part of them could just as well be regarded as the equipment necessary to complete the configuration of a new acquisition, given that only minimal peripherals have been included in the costing of that acquisition in this chapter.

9.1.6.2. A first model for VM/CMS

In this model VM/CMS is introduced on the existing computers using a system product that allows VM and MVS to coexist on the same machine. We need to start this new service in such a way that it does not interfere with the existing WYLBUR facilities. In order to do this we need some extra channels and main memory. Eight more channels cost 230 KSF, and eight megabytes of memory for the 3081 cost of the order of 320 KSF. This would allow a service with MVS and WYLBUR on the Siemens, together with MVS batch on one cpu of the 3081, whilst the other cpu runs a controllable mixture of MVS batch and VM/CMS. A Series 1 controller would cost 85 KSF, and permit 32 users to have full-screen CMS facilities via the existing Pericom terminals. Hence the first stage, which would be adequate for 1983, would cost of the order of 635 KSF. This should be more than enough to get VM/CMS installed, usable for system development during production (thus saving some production time), and permit all the existing TSO/WYLBUR users to convert to VM/CMS. At that point the longer term need could be reassessed.

Assuming that the assessment was positive, a second stage during 1984 would add one or two Series 1's, some disks and controllers, and would have to make a choice between another 8 MBytes of memory and a solid-state disk. The price of the latter depends very much on size, but a starting configuration of 22 MBytes would cost of the order of 450 KSF, somewhat more than the memory. Theoretically, is is probably always better to add memory rather than solid-state disks, but there is another factor to consider. The distribution of services between cpus, that we have outlined above, has one weakness, in that the Siemens 7880 machine has no microcode facilities. It is thus unable to run the above combination of MVS and VM on the same cpu anything like as efficiently as the 3081, and hence would lead to a fundamental asymmetry in services. A machine swap would lead to degradation of services. The Siemens (or ICL) M380 machine does have microcode, and would probably remove this problem. We have suggested that such a machine could be installed in 1984. If it is, then main memory is probably the better choice, but if we stick to the 7880 we should buy the solid-state disk, which at least could be swapped with the service if necessary. In summary, the second year of VM/CMS would again involve an investment of the order of 600 KSF.

At the end of this second stage, we could therefore have added between 64 and 96 simultaneous time-sharing users on top of the present 180 simultaneous WYLBUR users. Globally, this is not out of line with the numbers of access ports we have been adding over a similar period in the past to WYLBUR. We cannot tell whether the balance will work out correctly, whether the number of WYLBUR users converting to VM will balance the demand for WYLBUR from new users. It therefore might be necessary to add a small number of WYLBUR lines and go slower on VM. One could expect, however, that the Rutherford Lab. users, and those that have been at SLAC, would convert as soon as possible in order to use a system similar to the one they are used to.

A third stage to reach of the order of 150 simultaneous VM/CMS users would similarly cost of the order of 600 KSF. Further stages should continue as above, without the need for detailed elaboration here. Sooner or later a decision would have to be made about the replacement of a CMS batch service to replace the MVS batch. This has the advantage of homogeneity for the users, and the disadvantage of conversion.

In summary a VM/CMS service could be built up to a reasonable level in the next three years, at a modest expenditure of the order of 1.8 MSF. Much of this money represents extra equipment that is needed to complete the configuration of the new capacity as described in our model. If it should turn out after the first stage that some other system should be used on the IBM in the longer term, the memory and other equipment added would not be wasted. It would certainly be needed for the alternative system. This remark applies equally to the investment in VM/CMS for further years. If it isn't VM/CMS then a similar sum of money will have to go into an alternative in order to cope with the increased demand for general timesharing.

9.1.6.4. An alternate model for VM/CMS

The model above assumed that there would still be sufficient central capacity to introduce a VM/CMS evaluation as an additional task on the existing mainframes. The recent evolution of the load, in particular from UA1/2, has shown this to be a difficult route. If for no other reason, there will be tremendous psychological pressure against VM/CMS when the centre is heavily loaded. Hence we explain briefly an alternative route, which seems more realistic than the original model.

In the alternative model, a small machine like a 4341-9 and a small number of peripherals are acquired for about 500 KSF and used for the evaluation. Although a modest user service could be built up via progressive field-upgrades, it is more likely that a serious user service would start only at the time of a major capacity upgrade. One advantage of this route is that the 4341 would then become the basis of the file server machine.

One important point is becoming clear both in this model and in the way we have chosen to run WYLBUR since Christmas. It is desirable to separate, during the prime shift, the interactive computing and the batch wherever possible by running them in separate computers. It is a further illustration that the interactive computing is taking a significant fraction of the resources, and hence the cost.

In terms of overall money this model changes very little. However the individual budget lines are modified. The file server line is reduced by the purchase of the 4341 on the VM/CMS line in 1983. The major mainframe line, which previously contained an absolute minimum of peripherals, is increased by much of the extra equipment that we had labelled VM/CMS in our overall model. Since we consider the second model to be more realistic than the first, the cost table shows the model in this section.

9.1.7. Computer Centre Maintenance

The installation of more equipment in order to serve more users, provide more capacity and more services, inevitably implies more maintenance bills. The actual cost is extremely difficult to predict for the following reasons. Firstly, manufacturers tend to adjust their maintenance prices regularly for e.g. cost of living increases, (a situation which is embarrassing given the present total commitment of the computer centre budget, and the absence of an automatic allocation of a cost variation index). Secondly, it depends strongly on the type and age of the equipment. In general terms the maintenance on mainframes drops with new technology, e.g. the monthly maintenance of a 3081K in 1983 is 70% of that paid for the 168 in 1977, despite the factor 4.5 increase in power. On the other hand, if one doesn't change machines to take advantage of this, one ends up paying very highly for maintenance as in the case of the 7600. Peripheral maintenance costs tend to change much less than for mainframes, as one can imagine from the technology involved. Hence, the numbers in the cost table are calculated on the basis of a percentage of the capital cost of the new equipment.

9.1.8. Software Charges.

The computer centre will have to pay more than 700 KSF for software during 1983 for the main IBM and CDC services alone. It is now a well established policy of all computer manufacturers to separate software and hardware prices much more than in the past. Whilst the hardware costs for a given power drop regularly, the software charges rise with each new release. We already know that some of the announced system releases that we have not yet installed will cost very significantly more than the versions we are running today. For this reason we must foresee an increase of the order of 200 KSF at some stage, probably in 1985. Furthermore one must not forget that any additional mainframe would add of the order of 250 KSF for additional software licences. This is not foreseen in the model of this chapter, but could well become a reality, for example if we were to receive a machine from a LEP collaboration, which would decrease the probability of exchanging the 7880.

The trend to install new application packages for new developments like CAD, relational databases, etc. will mean that we have to pay more, and many of these packages are far from cheap. However we consider that application software should be included in the budget line for the application, and the line in the cost summary table refers only to the mainline computer centre operating systems, compilers, and support packages like HSM.

9.1.9. Personal Workstations

We have explained our belief that personal workstations, (which will of course be more often shared than personal!), will have a major impact on the way we compute during the decade, and thus must be taken very seriously. We have also stated that the time is not ripe to purchase large numbers of them today, however much people would like to have even today's offerings. We are in the process of purchasing 2 stations, and we should continue to purchase them in small numbers, for evaluation and learning, for the next year or so. Two of them are certainly not enough for today's demand, and one cannot expect people to continue to schedule their work regularly out of normal hours, if one wishes them to be productive. We are seeing new offerings in the market regularly, and with HP already in the market, we must expect the appearance of other major manufacturers, for example DEC and IBM. We must also expect a range of such devices, stretching from the size of the IBM personal computer, (not really a PWS), to something substantially more powerful than the colour screen APOLLO. As an illustration, APPLE have just announced their "Lisa" computer, which has a Motorola M68000 microprocessor like the APOLLO, a bit-map screen, a hard disk, and a mouse. The price is 10 thousand dollars for the Lisa, and said to be 2 to 5 thousand dollars for the unannounced cut-down "Macintosh" version. Since this promises to bridge the gap between the personal computer and the personal workstation, it could make a strong entry into the field, and also have a very interesting effect on pricing.

The problem is therefore is one of deciding when to move and to make a heavy investment. It is probably **not** 1983, but it could well be 1984. We will therefore anticipate, for the purposes of costing, a level expenditure of the order of 200/300 KSF during 1983, followed by a cost ramp, beginning in 1984, and reaching about 1 to 1.5 MSF per year in 1986/1987. It is hard to predict how fast these devices will fall in cost, and therefore to give sensible numbers. The stations we are now purchasing are fully equipped models at the upper end of the range, necessary in order to start serious work. However it is clear that, once the LAN and connections to the centre are in place, not all the stations will require disks. Nor will all the users require big screens, or even floating-point accelerators. Hence there is a second problem in guessing the configuration of an "average station".

The following numbers are pure guesswork, but let us suggest an average cost in late 83/84 of 50 KSF, and 35 KSF during 1985. By the end of the decade we might expect 10 KSF. With these numbers we could purchase 4 stations in 1983, 10 during 1984, 30 during 1985, 50 during 1986 and so on. So far we have been discussing only those stations which are funded centrally. If we are correct, we can expect to see more stations being purchased from other CERN budgets, when people become convinced of the justification, and, of course, PWSs will appear on site from other institutes that will, or already have, purchased them. There could well be 150 PWSs at CERN by 1987, and 500 by the end of the decade. However, even in the two months since these numbers were first estimated, there have been a number of price cuts that indicate that these costs are probably too high, e.g. the average station is already not far from 50 KSF. Hence it is quite possible that the number of stations in 1987 could be nearer 250 than 150.

9.1.10. Summary of costing

Table 2 shows a summary in 1983 figures of the budget lines explained and justified in the text of this chapter. The reader is invited to refer to the relevant section in each case. The base line is taken as the 9 MSF operating budget of the computer centre in 1983 which, as we have explained, represents the absolute minimum required to operate the existing equipment. This contains no reserve whatsoever for any purchase, and must be indexed in the future to cover increases in maintenance prices. Hence everything else in the table must be seen as an increase on this base number. The end total for the computer centre may be compared with the actual expenditure for the last three years, which in 1983 figures works out at an average annual spending of 14.3 MSF. This is 5.3 MSF above the 9 MSF base-line. In comparison with this figure, the WG3 model requires an expenditure in 1983 prices, which is 2.1 MSF in 1983, and very close to 5 MSF in each of the following 4 years. In other words, with the exception of this year, the total annual computer centre budget must remain very close to its level of the past 3 years.

1983 is a special case, where the major problem is to add enough peripherals to balance and exploit our existing configuration. We hope to delay any major new computer acquisition until 1984. Time will tell whether the delay is realistic. The pressure from the UA experiments, during a year when it is very inconvenient and inefficient for them to follow the 1/3:2/3 policy, may cause CERN to act earlier.

Table 2: Summary of Costs according to the Model in the text

	1983	1984	1985	1986	1987
Base operating budget of present computer centre at 1983 prices	9.0	9.0	9.0	9.0	9.0
Double Batch Capacity by 1987	0.0	2.6	2.6	3.0	3.0
Upgrade MSS to A3 model Later upgrade/replacement?	0.8	0.0	?	?	?
File Server + "Hyperchannel"	0.0	0.3	0.6	0.5	0.0
Printers, central plus RIOS replacement	0.2	0.6	0.3	?	?
General Peripherals to meet demand	0.4	0.75	0.75	0.75	0.75
VM/CMS TOTAL (note 1)	0.5	0.3	0.0	?	?
Extra maintenance at 3%/year of new central equipment	0.05	0.17	0.3	0.42	0.53
Increased software charges	0.0	0.25	0.45	0.45	0.45
TOTAL ADDITIONS COMPUTER CENTRE	1.95	4.97	5.00	5.12	4.73
TOTAL COMPUTER CENTRE	10.95	13.97	14.0	14.12	13.73
P.W.S.	0.2	0.5?	0.8?	1.0?	1.2?

Note 1. The costs for VM/CMS follow the second cost model in the text. A small machine is purchased for the evaluation phase, which is used afterwards for the file server.

Note 2. Not included in this table are microprocessor support VAX, CAD electronics, VAXs purchased out of experiment budgets even if used for general computing.

Note 3. All items are given in 1983 prices.

We know that at least one LEP collaboration may contribute a major computer to CERN. We are sceptical how much of this capacity will be available for general use. We realise that many LEP physicists consider the WG3 minimum increase of a factor 2 by 1987 as insufficient, and certainly not enough to cope with the simultaneous pressure from 4 experiments during the first year or so of LEP when, as UA1/2 are finding now, they will find it unrealistic to move the initial production and analysis outside CERN. Indeed the current reaction of some LEP groups goes in the direction of each experiment providing substantial **private** computing power, and installing it in the computer centre. Extrapolating to the limit, we would end up with the 1/3:2/3 policy being administered 1/3 in the computer centre on CERN purchased machines, and 2/3 in the computer centre on privately purchased machines. We believe that this desire for autonomy of the LEP experiments implies that CERN will still have to follow the WG3 model, and double the CERN-purchased computer power in order to provide for its 1/3 of the required capacity for LEP and for the rest of the programme. Time will tell, but the LEP collaborations are in need of an early indication of CERN's commitment in this area so that they can make their own plans.

9.2. MANPOWER

We have seen that, throughout the report, there has been a constant emphasis on the problems posed by the multiplicity of choices and systems on the market. Each of the choices has its supporters, and sound reasons to back it. The problems pose themselves in two ways. How can the users cope with having to learn multiple systems? How can the central staff support multiple systems? We have concluded that the answer is rather that it may not be the user who has to learn all the systems, he may in fact have to learn little more than he knows today, in that he may find at CERN the system he has at home, or that his work is confined to only one area with its own specialist system, or that the improved human interfaces that are coming will help him to work better. The problem is truly one for the support staff who **will have to support more** than they do today, because we have found very few realistic ways to restrict the number of options that are posed. Hence, support manpower is going to have to be somewhat larger than today.

9.2.1. User support staff

In the past 3 years we have doubled the number of users per week who submit batch jobs, and the number of WYLBUR users per week has grown even faster. The community of users has spread out across the laboratory. All our support services are geared in style, and in the education of the user support staff, towards helping the physics community and its particular style of working. That style is inappropriate for adequate support of the new users; engineers, secretaries, administrative staff etc. These people require packages more closely tailored to their needs than the physicists, and a better human interface. They request formal training courses, something which the physicists have tended to reject. The computer centre, and in particular the user support group, is just not set up or staffed to fulfil the needs of the non-physicists. Furthermore, they are more than fully occupied with what is still their major work, which is physics support. (Just to counter any false impression that may have been created by the above, the number of physicists using the computer has also increased significantly). This has been a growing problem over the past few years.

In order to resolve this, a number of things must happen. Firstly, the users must become more able to help themselves. This can be achieved partly through more modern systems such as those on workstations, or VAX/VMS, or interactive systems in general, but this is not enough. It also requires an improvement in the level of education of the users in computing, which means training. The people to give the courses do not necessarily have to be found within the user support group, nor even necessarily inside CERN. The management must accept the need for training, provide the funds to pay teachers if necessary, and accept that the people must be given the time off from their work to attend the courses.

Finally, it requires an increase in the number of people involved in supporting users, be they directly in the computer centre user support group or not. Many physics groups today receive good first line support from applications programmers within their groups. It is especially **outside** the physics community that such first-line local support is lacking. Hence staff within those areas must be trained to provide this support within their own communities or specialised application areas. Nonetheless new general services will have to be staffed directly within DD, and new posts will have to be found. One such area, of specific concern to LEP collaborations amongst others, is the lack of any formal user support (apart from the program library) for general service VAX/VMS machines.

9.2.2. Operational staff

As far as operational manpower for the centre is concerned, we can see few ways of reducing it. Tape mounting is a major problem, and one which is growing with time, despite the tremendous past success of the mass store in keeping the tape mounting load flat for several years. Today we are for long periods at the limit of the number of tape mounts that can be achieved, and we clearly lose some production capacity as a result of being unable to mount the tapes fast enough. Studies made of the ATL (Automatic Tape Library) device that is installed, for example, at Berkeley and Fermi Lab., reveal that it is a mixed blessing. The "juke-box" technology is clearly a problem for 100% availability, and the numbers of tapes it can hold would only be sufficient to make an impact on the tape mounting load if major efforts were made to organise the high traffic of experimental tapes which go in and out of the centre daily. It would also require a major capital investment greater than 1 MSF.

As another way of saving manpower, one could consider decreasing, (even to zero), the staff in the centre between say midnight and 06:00. This is certainly common, or even required by law, in other centres. However, other centres do not have our tape mounting load. We would certainly lose a sizeable fraction of our production capacity by such a measure, since the mass store is not big enough, by a large factor, to store all the tape-based data needed overnight. We would therefore have to invest faster than otherwise in new capacity to handle the load during normal hours. Furthermore, almost all of the essential systems management activities, such as backup and other file-base tasks are carried out at night. Apart from the fact that they often need human attention because of problems, these would have to be pushed partially into the prime shift, during which there are enough problems already. With the increasing complexity of installed systems, systems management tasks are a rapidly growing consumer of operational manpower.

An investment in better remote printers, and improved geographical distribution, could reduce the paper tearing load, and this may well come in the next years. No suitable product exists today. In any case the number of staff per shift involved in paper tearing is less than one per shift, and can never go to zero. As to the numbers of console operators, the trend is to increase the numbers of machines and consoles in the room, e.g. special purpose VAXs, so there is little chance for a significant reduction there. The number of operational staff involved in supporting traditional remote activities has already been reduced close to zero, by moving the work onto the users themselves. However, the effort in central support and operation of networks and data communication systems is a substantial amount and growing.

In summary, it will be difficult to reduce the operational staff over the next five years even if major capital investments are made, and even then the solutions are hardly ideal.

9.2.3. Systems staff

In terms of the final major manpower component of central services, namely systems support, the mere fact that the number of systems is increasing rather than decreasing will produce increased demands on the systems programmers. Indeed, people who are perhaps applications programmers today will also have to become their own systems programmers. Many distributed installations will have to be supported mainly by the user groups themselves, e.g. one finds already within physics groups with VAXs, that one physicist spends of the order of half his time on managing the system. A greater number of programmers within DD will have to take responsibility for specific systems or application packages, as has happened for CAD and the LEP database system.

As in all the sections on manpower in this report, WG3 has assumed an approach of making the most of the manpower available, stopping development on one system for a given time in order to develop a new one, and finding new staff from inside existing application areas. We have assumed that new manpower will be very hard to find in the difficult LEP construction years. We have tried at all stages to reduce the number of different systems in our model, but even before the general publication of this report we have complaints about areas we have cut out. Our conclusion is that it will be impossible to accomplish the proposed program of work without some modest increase in system staff, simply because we will have to support significantly more systems.

In summary, the major central systems support will have to come from the (slightly expanded) community of existing computer experts, largely, but not only, within DD. More programmers will have to become systems or application managers. These professionals will have to understand even more systems than in the past. On top of this, since there will not be enough computing staff to do all that is necessary, we shall see more professionals from other fields, for example, physicists and engineers, becoming trained as computer systems managers for their own particular application areas.

9.3. Concluding remarks on costs and manpower

It remains probable that there are still people within the laboratory who regret the infiltration of computing at CERN, and find it hard to accept that the computing budget and manpower involved cannot be reduced, indeed must be gently expanded when one adds the components in all Divisions. The answer to them is that computing has become, and will continue to grow as a vital element in the tools that are needed in this highly technical laboratory. Computers are essential if the professional physicists and engineers are to accomplish the programme of the next decade which represents a tremendous challenge. Computers will enable those people to be more productive than at present, and it is only through a strong programme in computing that we will be able to meet our goals. The final remark that must be made in answer to the above sentiments is as follows:

> "It is unworthy of excellent men to lose hours like slaves in the labour of calculation".
>
> Gottfried Wilhelm Leibniz (1646-1716)

Appendix A

MEMBERS OF WORKING GROUP 3

C.M.Jones (Chairman)	J.N.Gamble (Secretary)
R.K.Bock	R.Brun
J.M.Ferguson	S.Jarp
H.Kowalski	N.A.McCubbin
M.B.Metcalf	L.M.Robertson
S.Squarcia	

Of the members of the LEP Planning Group, who are de facto invited to the Working Groups, the following have attended WG3 regularly:

J.M.Gerard	M.G.N.Hine

The following people have been invited to specific meetings of WG3 to help us on specific points:

A.van Dam (Brown University)	
D.Lord	A.Hutton
P.Kunz (SLAC)	O.Martin
P.Watkins (Birmingham University)	

In addition, we are grateful for the written contributions of the following people who, whilst being too far away to attend WG3 meetings, have nonetheless followed our work with interest:

R.J.Cashmore	R.P.Mount
H.Newman	

We thank all the other people, too numerous to mention, for reading this report at various stages and for contributing helpful comments. Finally we must express our appreciation of the cheerful help of J.Megies in organising, copying and distributing the mountain of paper that WG3 has produced.

Appendix B

GLOSSARY

This glossary assumes that the reader is connected with the world of HEP! (High energy particle physics), and that his major problem will be in digesting the plethora of abbreviations that are found in the world of computing.

ADA: Language initiated by the United States Department of Defense as a common language for programming real-time systems. It is important because the DoD is the largest software consumer on earth.

Apollo: A personal work station (PWS) which combines a fairly powerful processor (roughly equivalent to a VAX-750), a high quality bit-mapped display, a local area network, a good Fortran 77 compiler, and a modern operating system.

Apple, Lisa, McIntosh: Apple is probably the most popular home computer around today. Lisa is a much more powerful and sophisticated computer recently announced by the Apple Corporation which includes a high resolution bit-mapped display and a mouse. McIntosh is a cheap version of Lisa.

Bit-map: Literally a map of bits. A screen with a bit map can be addressed at the individual bit level, giving much greater resolution and flexibility than conventional terminal screens with a fixed set of characters at fixed positions on the screen.

CAD,CAM,CAE: Computer Aided Design, Computer Aided Manufacture, Computer Aided Engineering.

CDC 7600: A very powerful computer first delivered in 1971, and for long the fastest general purpose computer available commercially. Despite its discrete technology the 7600 at CERN is still running 11 years after installation.

CMS: See VM/CMS.

CPU: Central Processor Unit. The heart of the computer.

Cyber 205: A powerful array processor sold by CDC, and developed from their Star 100 system.

Cyber 835: A medium sized computer sold by CDC which is capable of operating in 60-bit (6600 compatible) or 64-bit mode. The latter capability cannot yet be exploited since the 64-bit operating system is not available.

Cyber 875: A computer which will be delivered by CDC in 1983 which is roughly equivalent to a dual-cpu 7600.

CRAY 1, 1S, 1M, X-MP: The CRAY 1 is a powerful array processor designed by Seymour Cray who was responsible for the CDC 6600 and 7600 machines. The 1S, 1M and X-MP are more powerful versions using the same basic architecture as the CRAY 1.

DBMS: Data Base Management System.

DST: Data Summary Tape in high energy physics. Such tapes are written at various stages of the processing of the raw data tapes that are written during data-acquisition.

ECFA: European Committee for Future Accelerators.

ELECTRIC: A conversational editor, file management and job entry system developed and used at Rutherford Appleton Laboratory on large IBM computers.

Emulator: The name given to a family of microprogrammed devices that emulate the machine instructions obeyed by a larger computer. In high energy physics this was first used to execute the instructions produced by the IBM FORTRAN compiler. The basic idea was that many such devices could be processing individual events simultaneously, thus achieving power through weight of numbers. Many variations now exist, e.g. 168E, 3081E, 370E.

ERASME: CERN system combining in one phase the scanning, measurement, pattern recognition, geometric reconstruction and event rescue of bubble chamber events on film.

Ethernet: A local area network technique originally developed by the Xerox Corporation, for which many computer manufacturers have now announced support.

EXEC Files: A file containing a program written in a rudimentary language which is interpreted by Wylbur. Similar facilities are available with other conversational or interactive systems, e.g. CLISTs under TSO, Command Files under VMS.

FASTBUS: A specification for a very high speed bus developed for use with data acquisition electronics.

File Server: A computer dedicated to maintaining a file base on behalf of other computers.

HTV: An interactive version of the HBOOK/HPLOT histogramming package, which enables a terminal user to generate dynamically different graphical representations of his data.

HSM: Hierarchical Storage Manager. An IBM product which organises the datasets on a system running the MVS operating system, automatically moving files between real disk and mass storage according to file usage and available disk space.

ICL Atlas 10: The Fujitsu M-380 as sold by ICL.

Index: A data communications switching system which connects terminals to computers, permitting the terminal user to select dynamically the computer to which he wishes to be connected.

INFOL: A database system developed at the University of Geneva and CERN.

INTERCOM: A CDC interactive system running under the NOS/BE operator system.

ISO: The International Standards Organisation.

JES, JES2. Job Entry System. In IBM terminology this is that part of the operating system which handles batch jobs, printing. JES2 is the JES used at CERN on the IBM systems under MVS.

LAN: Local Area Network. A communications network which is geographically limited. Usually it is of relatively high speed (at least 10 megabits per second). Common LAN techniques are Ethernet, Cambridge (slotted) Ring and Token Passing Ring.

LCN: Loosely Coupled Network. A very high speed (50 megabit per second) local area network developed by CDC for interconnecting mainframes and disk systems.

LEP: Large Electron Positron Collider. A 27km circumference accelerator being built at CERN. The present working group's raison d'etre.

Lisa: See Apple.

McIntosh: See Apple.

Mass Store: Computer storage which has substantially greater capacity than that available on most systems. Thus whereas ten years ago disks were often referred to as mass storage, today the term usually refers to systems with at least 10^{12} bits accessible without operator intervention.

MIPS: Million Instructions per Second. A rough measure of computer power.

Mouse: A pointing device developed at Xerox Palo Alto Research Centre as a means of indicating cursor movement on a terminal. It has a small ball built into its base which rotates as the operator moves the mouse across a surface (the desk top). The mouse senses the direction and speed of the rotation and these are used to define the movement of the cursor across the screen. The mouse usually also has some buttons which allow the operator to signal additional functions.

MSS: Mass Storage System. This generally refers to the IBM 3850, a cartridge tape system.

MVS: Multiple Virtual Storage. One of IBM's mainline operating systems for the 30xx and 4341 computers, providing very high functionality and good batch processing capabilities, but with an antediluvian control language.

MVT: A former IBM batch operating system, designed for the 360 architecture.

M380: A Fujitsu computer of power greater than 10 times a 168, in double processor version.

NAG: The National Algorithms Group. A British organisation which maintains libraries of standard mathematical algorithms on a wide variety of computers and operating systems.

NORD: The name given to computers manufactured by Norsk Data, a Norwegian computer company. Models vary from the Nord-100 to the Nord-500.

NOS: Network Operating System. CDC's mainline operating system, geared towards interactive use.

NOS/BE: Network Operating System/Batch Environment. A CDC operating system developed from the former SCOPE 3 system, oriented towards batch processing. It is no longer being developed by CDC.

Pascal: A block structured programming language developed by N.Wirth of the ETH, Zurich. It is widely used as a teaching language, and has recently been standardised by the ISO.

PATCHY: A source code maintenance program developed at CERN.

PERQ: A personal work station manufactured by ICL which combines a powerful cpu with a high quality bit-mapped display. It is capable of very impressive graphics.

Print Server: A computer dedicated to the function of printing files on behalf of other computers.

Puck: A pointing device used in conjunction with a graphics tablet which not only indicates position, but also has one or more buttons which can be used to select functions.

PWS: Personal Work Station. A computer sufficiently inexpensive to be dedicated to a single user. We distinguish a Personal Work Station from a Personal Computer (PC) by its greater power.

Rasterops: The name given to a set of instructions often included in personal work stations with bit-mapped displays to enable the display memory to be updated efficiently.

SCRIPT: A document formatting program developed by the University of Waterloo which runs on IBM computers. This report was formatted by Script.

SPF: Structured Programming Facility. An IBM product which combines an editor and a set of (mainly file manipulation) utilities with a tree structured menu system and full screen support of IBM 327x terminals.

SQL: The non-procedural relational database query language developed by IBM and now used by several relational database systems.

SVS: Single Virtual Storage. An IBM operating system developed from MVT which took advantage of the virtual address capabilities of the 370 range of computers.

Star: The Xerox Star is a personal work station which offers a particularly advanced operator interface, using 'icons' (pictures of files, in-trays, etc.) instead of words to guide the user.

TSO: Time Sharing Option. An IBM interactive system running under the MVS operating system.

UA1/UA2: Underground Area experiments 1 and 2. The two major experiments installed in the underground halls of the CERN SPS when used as a proton-antiproton collider. They have recently produced very strong evidence for the discovery of the long-sought W particles.

UNIX: A popular, portable interactive system developed at Bell Labs for small DEC machines, but now available on a wide range of computers. UNIX (or UNIX look-alikes) is fast becoming the standard operating system for 16 bit microcomputers and workstations.

VAX: The name given to a range of 32-bit computers manufactured by DEC.

VM/CMS: Virtual Machine / Conversation Monitor System. An IBM operating system running on 43xx and 30xx series machines, providing efficient support of large numbers of interactive users.

VMS: The virtual memory operating system for DEC's VAX computers.

Wylbur: A conversational editor, file management and job entry system running under the MVS operating system on IBM machines, programmable only through an interpretive language facility (EXEC files). Originally developed at Stanford University it is now widely used in university and laboratories.

168: An IBM computer, in its time (1976) the top of the IBM range.

168E: An emulator built at SLAC and CERN, which aimed to be one third as powerful as a real 168. The 168E can execute only a subset of the 168 instruction set, but sufficient to enable it to execute code produced by the Fortran compiler.

3032: An IBM computer similar in power to a 168.

3081: The most powerful computers currently manufactured by IBM, in 4 cpu version approaching the power of 10 168s.

3081E: An emulator being developed at SLAC and CERN.

370E: Another emulator of the IBM instruction set.

7880: A Fujitsu computer as sold by Siemens, of power around 2.3 times that of a 168.

7890: The Fujitsu M-380 as sold by Siemens.

Appendix C

INDEX OF WORKING PAPERS OF WG3

The following is an index, sorted by reference number, of the WYLBUR library (PDS) CM.WG3.PRINTER in which all the working group 3 papers are stored. Most papers are in a form suitable for direct printing, for example via the LISt OFFline UNNumbered CC command. Unfortunately any graphs or figures are not in the text. Similarly, a small number of papers are not in computer form, (although they exist in the index), and these may be obtained via Mrs J.Megies, Computer Science Library, DD Division, CERN (internal telephone 2379).

AGENDA1	WG3/cmj/82/1,	Agenda for the first meeting
TERMS	WG3/cmj/82/2,	Terms of reference and first list of topics
MIN1	WG3/jng/82/3,	Minutes of the 1st meeting 4 June 1982
SERVICES	WG3/cmj/82/4,	List of Services
MIN2	WG3/jng/82/5,	Minutes of the 2nd meeting 11 June 1982
MIN3	WG3/jng/82/6,	Minutes of the 3rd meeting 25 June 1982
COMSCEN	WG3/mgnh/82/7,	Scenario of future communications services
MCCSCEN	WG3/NMcC/82/8,	Possible computing scenario for LEP
CENSCEN	WG3/jmf/lg/82/9,	Brief consids. Central siting of computers.
JNGC90	WG3/jng/82/10,	Scenario of distributed computing for CERN
GENSER	WG3/mgnh/82/11,	Remarks on future communication services
ECFAMEM	WG3/ecfa/82/12,	ECFA computing for LEP memo
ECFALIST	WG3/ecfa/82/13,	ECFA members address list
LEPORIG	WG3/lep/82/14,	The original Dobinson/Zacharov report
ECFANET	WG3/ecfa/82/15,	ECFA wg5 on Wide Area Communications
RATEGRPH	WG3/nm,ss/82/16,	Graphs of event and tape rates
MIN4	WG3/svj/82/17,	Minutes of the 4th Meeting 2 July 1982
MIN5	WG3/nam/82/18,	Minutes of the 5th meeting 9 July 1982
SCENBRUN	WG3/rb/82/19,	Scenario for future services
MIN6	WG3/jng/82/20,	Minutes of the 6th Meeting 22 July 1982
MOUNTCOM	WG3/rpw/82/21,	Comments of R.Mount on LEP Computing needs
MIN7	WG3/lmr/82/22,	Minutes of the 7th Meeting 6th August 1982
MIN8REV	WG3/nam/82/23,	rev. Minutes of the 8th Meeting 13 August 1982

POLBOCK	WG3/rkb/82/24,	Concepts for a computer policy
WG3MOD	WG3/cmj/82/25,	A Model for LEP Computing for discussion
MIN9	WG3/jng/82/26,	Minutes of the 9th meeting 20th August 1982
MIN10	WG3/jng/82/27,	Minutes of the 10th meeting 3rd September 1982
APOLTRIP	WG3/br/82/28,	Trip report from APOLLO
MIN11	WG3/jng/82/29,	Minutes of the 11th meeting 16th September 1982
INTREPR4	WG3/cmj/82/30,	Interim Report, final revision 4.
MIN12	WG3/jng/82/31,	Minutes of the 12th meeting 30th September 1982
COSWYL	WG3/cba/82/32,	Costs of the WYLBUR Service
JAMBOREE	WG3/jng/82/33,	Summary of the First Open Meeting
CODEMAN	WG3/hg/82/34,	Future CERN source code manager
ELECMAIL	WG3/lmr/82/35,	Draft on Electronic Mail
MOUNT2	WG3/rm/82/36,	Comments by R.Mount after the Jamboree
MIN13	WG3/jng/82/37,	Minutes of the 13th meeting 22nd October 1982
MIN14	WG3/jng/82/38,	Minutes of the 14th meeting 29th October 1982
HIGHLL	WG3/mm/82/39,	High-level Languages on CERN Mainframes
CSNET	WG3/mgnh/82/40,	Report for information on CSNET
VOICEMAI	WG3/mgnh/82/41,	Note on Voice Mail.
MIN15	WG3/jng/82/42,	Minutes of the 15th meeting 5th November 1982
TSOVVM	WG3/jg,cj/82/43,	TSO vrs VM discussion.
MIN16	WG3/jng/82/44,	Minutes of the 16th meeting 11th November 1982
LMRINT	WG3/lmr/82/45,	Interactive Computing
MIN17	WG3/jng/82/46,	Minutes of the 17th meeting 18th November 1982
SOFTMAN	WG3/ms/82/47,	A software management system
CADELEC	WG3/mfl/82/48,	revised. CAD Electronics note to DD/JPC
RPM3	WG3/rpm/82/49,	R.Mounts comments November 1982
NEWMAN	WG3/hn/82/50,	Comments of H.Newman
REFREP	WG3/hg/82/51,	Referees Report on Interim report revision 2
MIN18	WG3/jng/82/52,	Minutes of the 18th meeting 26th November 1982
FINREPR3	WG3/cmj/83/53,	Final Report revision 5. 17th March 1983
JAMB2	WG3/jng/82/54,	Summary of the 2nd Jamboree
LIBS	WG3/hrr/82/55,	H.Renshall's note on Program Libraries
MIN19	WG3/jng/82/56,	Minutes of the 19th meeting 13th December 1982
MIN20	WG3/jng/83/57,	Minutes of the 20th meeting 20th January 1983

REFERENCES

[1] V.Zacharov, "Invariant Coding and Parallelism in Data Processing", CERN Data Handling Division, DD/82/1, 1982.

[2] Workshop on Software in High-Energy Physics, CERN Yellow Report, CERN 82-12

[3] Workshop on Software in HEP, Conclusions, ECFA/82/62

[4] M.Metcalf, A Proposal to link WYLBUR to PATCHY, DD/CUAC/78/372

[5] P.Hazel, Resource Allocation and Job Scheduling, University of Cambridge, Computing Laboratory, Technical Report No. 13, 1980.

[6] G.Waldbaum, IBM T.J.Watson Research Center, Yorktown Heights, N.Y., Proceedings of the SEAS Anniversary Meeting, 1981, held in Nice, Volume 2, pages 557/568.

[7] Joint Training Board, CERN TRAINING PLAN 1982/1983, December 20, 1982.

CONTENTS

TABLES

FIGURES

APPENDIX E
TO THE LEP COMPUTING PLANNING GROUP REPORT

PGI/cr

10 December, 1982

Final Report of the Steering Committee for Microprocessor Standardization

Introduction

1. The Interim Report of the Steering Committee for Microprocessor Standardization was issued at the beginning of August 1982. It contained a set of firm recommendations and asked for further work in some areas before other recommendations could be made. E. Gabathuler agreed on the continuation of the work and set up an extended Steering Committee, to complete the study with the help of two subgroups (see Appendix I for Committee and Subgroups membership and terms of reference).

2. The recommendations that needed no further investigation are : (the Interim Report should be consulted for details).

   - A Microprocessor Support Group should be set up in DD Division, to provide support for the chosen standards. It may be necessary to involve other divisions in the hardware support.

   - An interactive development service should be provided on a dedicated central 32-bit minicomputer.

   - The support should be concentrated on the 6809 and on the 68000 family of microprocessors.

   - Where an expensive development system is justified the HP 64000 is recommended.

   - The language for microprocessors to be fully supported is an extension of ISO PASCAL. FORTRAN should also be supported, subject to availability of a compiler.

   - For simple 8-bit applications the G-64 bus is recommended.

3. The topics needing further technical investigations before a recommendation could be formulated were:

   - a closer definition of the central interactive facility in relation with the user's needs and the availability of software. Temporary measures to cover gaps in the short term.
   - The analyses of the Preliminary Inquiry and choice of software.
   - The choice of the 16-bit bus and crate system.
   - A definition of the medium cost development systems based on the chosen bus and crate.

   In addition manpower, money, timescale and organization of the support were to be assessed and defined.

4. After analysis of the available implementations of Pascal, the LEP control system selected Modula 2 as system programming language.

5. Based on the work of the two subgroups reported in Appendix II and Appendix III and on communications from the ECFA subgroup, chaired by L.O. Hertzberger, and the LEP Computing Working Group I, chaired by R. Dobinson, the Steering Committee agreed on the following recommendations.

Recommendations

5.1 Central interactive facility. It should consist of a VAX-11-780 running UNIX, located in the computer centre, with connections to both INDEX and CERNET. CERNET concentrators at the user's end should be provided as a cheap means of supplying many users with down-line loading facilities of adequate speed.

5.2 Workstations. The possibility to offload the central facility in the long term should be investigated. Workstations based on the recommended 16-bit microprocessor, bus and crate should be tried out.

5.3 General software. A user-friendly centralized filing system and a source code management system should be provided centrally for both 8-bit and 16-bit users as a means of promoting standardization by encouraging sharing of code and procedures. Libraries of general usefulness in microprocessor applications should be centrally supported.

5.4. Portability. The cross software should in principle be portable, and where software is obtained from an external source CERN should ensure, as far as possible, that it can be distributed on reasonable commercial terms to collaborating institutes.

5.5 Cross Software for the 6809. The inquiry shows no PASCAL cross compiler available for the 6809 before the second half of 1983. Until that date, what is available centrally to 6809 users is a cross-assembler only. As a way to standardization in the short term, it is recommended that the 7500 series development system based on G-64 and running Omegasoft Pascal under the FLEX operating system be supported with connections to a centralized filing scheme.

5.6 Cross Software for the 68000. It is recommended that the Siemens Pascal cross compiler, now operational on the IBM and VAX/VMS, should be retained as the best choice. None of the firms contacted could offer a FORTRAN compiler for the 68000 at present. It is felt that a FORTRAN cross compiler should be available as soon as the central VAX support becomes operational; therefore it should be commissioned without delay.

5.7 Real-time Kernel. The RMS68K real-time kernel from Motorola is recommended for the 68000. Subject to evaluation currently in progress, OS-9 from Microware is a candidate for the 6809. However standardization efforts for real-time kernels are under way in different bodies. CERN should follow this effort and be prepared to switch if a suitable standard product emerges.

5.8 Choice of bus and crate systems. The projection of the use of modules fitting the standard buses over the next 5 years at CERN indicates approx. 2000 crates and 37000 modules of which 60% are "standard" in the sense that they are not specific to a project. Of the standard modules, the 8-bit range covers 17000 modules, the 16-bit range 5500 modules.

As already stated in the Interim Report, the G-64 crate and bus is retained as the 8-bit standard.

For the 16-bit system the VME bus and crate is recommended. A clear definition of the G-64 and VME options supported by CERN shall be made known.

It should be realized that the market situation is far from ideal for both G-64 and VME when compared to the requirements of the CERN purchasing policy. G-64 modules are available only from the CERN Host States; VME modules predominantly from the United States, at least for the time being. In this situation a way of promoting multiple sourcing from a number of firms spread over all CERN Member States is to ask for equipment built to CERN specification. It is recommended that purchasing of the standard modules be centralized in order to enforce a balanced policy. The standard modules should be made available and serviced through a central organization.

5.9 Medium cost development systems. Microprocessor development systems using G-64 or VME should be based on hardware and software building blocks (such as microprocessor, memory and communication boards, diagnostic and debugging software) defined and supported by the Microprocessor Support Group.

5.10 Integration. A central support team is required to ensure proper integration of cross software, real-time kernel, communication and application software with the medium cost development system.

Cross software installed on the central UNIX system will extend to local workstations running UNIX. When work stations built in the VME standard become available it will be possible to integrate the work station and the VME medium cost development system. The work station evolution should therefore be closely followed by the central support team.

The present ISR 7500 system provides a combined work station and development system for G-64. When 6809 cross software becomes available under UNIX it will be possible to use the central support machine or VME-based work stations for G-64 developments.

Timescale

(Assuming manpower and money can be assigned to this project without delay.)

6. The central interactive VAX running UNIX, if ordered at the beginning of 1983 and installed within 4 months, could be operational by the middle of 1983, with INDEX and CERNET connections.

7. At present, microprocessor cross software other than specified in this report is supported on the central computers. (See Appendix IV, for details.)

8. Work on MODULA 2 is in progress in DD aimed at providing a complete cross system for a number of target machines on the IBM (see Appendix V for a description of the aims and the expected timescale).

9. A PASCAL cross compiler suite for the 6809, if available for purchase in the second half of 1983, could be ready for CERN users on the central VAX by the end of 1983.

10. The SIEMENS PASCAL cross compiler suite for the 68000 is available now on the IBM and on a VAX under VMS. It will be operational on the central VAX in the 3rd quarter of 1983, assuming the installation of the VAX goes as in paragraph 6.

11. Hardware to the G-64 standard is available now and a CERNwide support could become effective early in 1983; the same applies to the 7500 development system. The availability of VME hardware is less immediate and will require a market investigation in the first quarter of 1983, as well as extensive experimentation. It is however reasonable to assume that medium cost development systems based on VME (see 5.9) will be available in the 3rd quarter of 1983 at the same time as the cross software on the central VAX. The CERN-wide support on VME could therefore start to be effective by the end of 1983.

Budget and Manpower

12. Two situations will have different budget and manpower requirements:

- The initial period, lasting of the order of two years, during which the purchase of hardware and software will be required to set up the service. This period will also require a great effort in providing and adapting software and in surveying the hardware market and promoting collaboration with industry.

- During the following 3 or 4 years innovations will be required by the technological evolution, but a large share of the activity will be software maintenance and improvements, advice to users and hardware purchase, acceptance testing and repairs.

13. The budget required for the first year is 1.6 MSF, subdivided as:

VAX purchase	.8 MSF
Software purchase (UNIX, cross compilers, etc.)	.3 MSF
Other hardware (prototypes or commercial for evaluation, test gear, essential developments)	.5 MSF

For the following years a running budget of .6 MSF should be foreseen subdivided as:

Computer maintenance	.1 MSF
Software maintenance and additions	.1 MSF
Hardware purchase	.4 MSF

It goes without saying that all hardware purchased for the users groups will be charged to their budgets.

14. Setting up the central service will require an initial manpower investment of 20 people, half hardware and half software specialists with a good overlap. As the variety and number of modules increase, additional staff must be added for hardware testing, acceptance and repairs. In addition an infrastructure is required to take care of inventory, accounting, shipping, receiving, etc.

These manpower requirements are based on the assumptions that in most cases software will be purchased and well documented hardware modules will be designed outside the central service. However some of the inevitable software and hardware development done at CERN should be within the central service to ensure a high level of expertise.

15. Support is provided at present in DD and ISR to microprocessor users, primarily in the DD-SW Group for languages and in the ISR-CO Group for G-64 and the associated 7500 development system. Moreover the DD-OC Group is shifting manpower progressively from minicomputer to microprocessor support. As VME will find applications both in accelerator controls and LEP experiments, the involvement of the EP Test and Instrumentation Group appears to be unavoidable, although this additional load does not match the available manpower.

Organization

16. The four areas mentioned in 15, represent the principal nuclei where knowledge and experience are concentrated: the central microprocessor support should be established by taking advantage of these assets.

It is proposed that an interdivisional project should be set up, bringing together people from the four areas under a project leader; additional staff will be needed to meet the requirements outlined in 14. In view of the close connections between hardware and software, geographical unification should be encouraged. The project should aim at setting up the microprocessor support service following the recommendations of this report. The project should have the duration of two years: thereafter the situation should be reviewed by the Management on the basis of achievements, users' needs and technological evolution.

17. Information related to organization, timescale, product availability, expected policy changes etc. shall be made available by the project through a regular bulletin (Microprocessor Newsletters) open to contributions and suggestions by the users of the service. In addition, a users forum shall be convened at regular intervals, to review the quality of the service and suggest corrections to the line followed by the project.

18. The CERN Technical Training should be given the resources to adapt its courses in the microprocessor field to the recommendations contained in this report.