EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH

CERN-LEP/DI/83-20

THE CONTROL SYSTEM FOR LEP

by M.C. Crowley-Milling

Presented at the 1983 Particle Accelerator Conference,
21-23 March, 1983, Santa Fe, New Mexico

Geneva, Switzerland
April 1983

THE CONTROL SYSTEM FOR LEP

M.C. Crowley-Milling
European Organization for Nuclear Research (CERN)
1211 Geneva 23
Switzerland

Introduction

Four years ago, during the design study phase of LEP, a paper¹ was presented at this conference, giving the then current ideas about the control of LEP, and proposing some developments and investigations that should be carried out. LEP is now an approved project², and the situation has been reassessed. Many of the statements made in that paper are still true, but others need to be modified, in view of subsequent developments and the change from the specialized injector of the original scheme to the use of the PS/SPS complex for this purpose. This was then followed by the decision to couple the LEP control system to that of the SPS; following the same principles³, and to make one group responsible for the two systems.

There are three main ways in which the requirements for the LEP control system differ from systems that have been built in the past. These are the great distances involved with a ring 27 km in circumference, the introduction of microprocessors into the equipment itself and the long timescale at a time when progress in electronics is rapid, with the cost of hardware falling every year, without any corresponding fall in the cost of specialized software.

In view of the latter, there are many detailed decisions that cannot be taken at the moment, but the basic principles have been laid down, and a number of developments and investigations have been started. This paper firstly discusses the effects of the introduction of microprocessors, followed by the basic principles of the LEP control system and a description of the work in progress.

The Introduction of Microprocessors into the Equipment

One of the biggest revolutions in control system design has been caused by the introduction of microprocessors as an integral part of the equipment to be controlled, taking over the duties of sequencing, local surveillance, local servo-loops, function generation and testing, thus reducing considerably the load on the control system. There are other advantages as well, as we can see by taking as an example a magnet power supply unit. Previously, such a power supply would need a number of control actions and status acquisitions to be provided by the interface equipment, as well as connections to a multiplexed analogue measurement system, and, if the magnet had to be ramped, there would be the need for an external function generator. Because of the cost of interfacing, the possibilities of remote fault-finding were rather limited, and all the external connections to the power supply were potential sources for interference pick-up. With microprocessors, the situation changes considerably. The large-scale integration has also brought down the price of precision analogue-to-digital converters so that a separate one can be fitted to each power supply. Microprocessors can be programmed to act as function generators, operating on a stored table of values previously loaded down from the control system, using the timing system for synchronization. Diagnostic and test programs can be run when something goes wrong, to provide the operators with full information as to what to do. A command interpreter can be provided so that the interaction with the control

system can be in the form of simple messages, instead of a succession of command/response actions. All this can be provided within the power supply unit itself, with just the connection for messages and timing as links to the outside world, reducing the cost and possibilities of interference pick-up.

Even when the equipment forming a system is more diffuse than the example used above, a number of items, such as a water pump and its associated valves and sensors, can be controlled as a unit by a microprocessor, rather than having to interface the individual items to a control computer.

The Control Computer Network

There are several different possible configurations for a distributed computer control system. In many cases, a large central computer is used to service a number of consoles, to run the main control programs and also to maintain the data base. The equipment containing microprocessors can either be connected directly to the main computer via a local area network of some kind, forming a two-level hierarchy, or groups of equipments belonging to one subsystem or in one geographical area can be connected to small computers, which in turn are controlled by the main computer, forming a three-level hierarchy.

For LEP, we propose to use a three-level hierarchy, but with the central computer replaced by a network of smaller computers, between which the duties of a central computer will be divided, in a similar manner to the SPS system. This results in a situation shown in simplified form in Fig. 1, where the various computers, some of them process computers directly controlling equipment and others driving consoles or providing specialized services, are connected together by a LEP-wide upper-level network, and each of the process computers is connected to its equipments by one or more lower-level local area networks.

Fig. 1 : Simplified system layout for LEP

Since the SPS/PS complex will be used to provide the electrons and positrons for injection into LEP, it has been decided that LEP will be controlled from an extension to the SPS control room, as having a combined control room will simplify the coordination and should help to minimize the operations personnel required. Therefore one of the requirements is that the upper level network should be easily interconnectable with that of the SPS. The SPS uses a multi-star packet-switching network that was designed before standards for such networks had been agreed. The present proposal is to use a similar multi-star network for LEP, using the same higher level protocols, but changing the lower levels to take advantage of standards such as HDLC. There does not seem to be a complete standard network available which will cover the special requirements of a control system, such as a guaranteed maximum response time for short messages, and also cover the area required, and so it seems likely that we will have to continue with our in-house development. We are meanwhile keeping watch on outside developments, such as the IBM token ring, to see whether the performance and time scale can meet our requirements.

It might be thought that a ring network would be more suitable for the control of a circular machine, but the multiple-star configuration also fits in quite well, as can be seen from Fig. 2, which shows a simplified layout for the proposed LEP network.

Fig. 2 : Simplified layout of proposed LEP network

We are considering two possible multi-drop highways for the lower level networks, one an in-house development, based on multi-drop RS-422 CSMA/CD, using HDLC frames, and the military standard multidrop system MIL-STD-1553B. The latter was developed for aircraft control, but its use is now spreading to process control applications. Since most of the equipment will be able to exchange messages with the process computers, the decision has been taken that all communication on these networks should be in message form, and special arrangements should be made for the connection of simple equipment that cannot interpret messages.

It has also been decided that these messages should be in the form of ASCII characters. This has the advantage that the messages can be in a printable and easily recognizable form, and that the equipment can be operated for test purposes from a very simple terminal.

The decision to limit the communication on the multidrop highway to ASCII messages means that special steps have to be taken to connect to any equipment that does not contain a microprocessor. For example, if it is wished to use some CAMAC modules, because they are particularly suited to a certain job, it will be necessary to provide a CAMAC crate controller which incorporates a microprocessor, to take the messages and turn them into the CNAF command/response mode required by the CAMAC protocol. Similarly, adaptors will be needed to control test equipment which is already interfaced to the GP-IB (IEEE 488) standard, which is becoming quite popular. Consideration is being given to standardizing all messages on the multidrop highway to follow the Tektronix codes and formats for the GP-IB.

The Computers

One of the difficulties experienced with computer control systems is that the life of an accelerator or storage ring is normally longer than the life of a computer. At the PS, at Fermilab and at Los Alamos, amongst others, the original computer(s) installed have become obsolete, are expensive to maintain, and in some cases the original manufacturer has given up making computers, or even ceased to exist. Even when the computer, or a more modern replacement, is still supported by the manufacturer, sometimes the old operating system is not and the new operating system requires extensive modifications to the existing software.

With a system like that of the SPS, using a central library, the process computers (those connected to the equipment) are fairly simple, using a real-time executive, and they do not require all the facilities of a comprehensive operating system. It was suggested some time ago that such a computer, where the processor has to be shared between a number of concurrent tasks, could be replaced by an assembly of microprocessors, each dealing with a single task. The executive would then be reduced to little more than a scheduler and arbiter. Compared with a minicomputer, such an assembly should be cheaper initially, but the main advantage would be on maintenance, which could be carried out by replacing a complete microprocessor board, rather than relying on an expensive, and sometimes unsatisfactory, maintenance service from the computer manufacturer. Also, additional requirements could be met by plugging in an extra module, or replacing one by a later, more powerful version.

This idea of using an assembly of many microprocessors, each performing a single task, has only been worth considering since the price of memory has fallen so low, because, to preserve the basic simplicity of the concept, each microprocessor should have enough memory to carry out its task, and should communicate with the others by passing messages, rather than using shared memory. Critics of this proposal point out that very soon powerful microprocessors will be available with a comprehensive multi-tasking operating system included on the chip, which will be able to replace the present day mini-computer directly, without going to the complication of the multi-microprocessor assembly. This may be true for such things as personal computers, scientific work stations, etc., but the situation is different for a real-time process control application, where the simplicity of the software, the ability to debug and test the different tasks independently, and the possibility of tailoring the hardware configuration to fit the requirements, make the multiprocessor assembly sufficiently attractive that it is proposed to use this scheme for the process computers for LEP. It is also interesting to note that Texas Instruments have based

their latest designs on the principle of performing specialized functions in separate processors. They call it "Function to Function" architecture. An experimental set-up has been made at CERN using a number of CAMAC auxiliary crate controllers to replace the majority of the tasks performed by a NORD-10 minicomputer in the SPS system, to see what difficulties might be experienced, before building prototypes suitable for use with the LEP system⁴.

CAMAC has its limitations as a basis for such a system, as it was not designed for multi-master operation. There are a multitude of card, crate and bus systems for microprocessor use in existence, mostly designed by microprocessor manufacturers to suit their particular components. Although some of these have become de facto standards, and have been recognized as such by the IEEE, this institution set up a Working Party to try to propose a standard for a comprehensive, microcomputer-independent, bus system, to which it gave the project number P896. However, there have been difficulties in reaching agreement over some of the details, and little progress seems to have been made recently.

The European nuclear electronics standardization committee ESONE has also been working for the last two years towards a standard bus protocol which has been given the name E3S. At the General Assembly this year, the ESONE Committee is expected to approve the specification for this protocol, which can be applied to several crate and bus systems in which multiple-mastership of the bus is possible. For the LEP multi-microprocessor assemblies, it is proposed to use the E3S protocol with the VME crate and bus system. More details of this proposal are given elsewhere^{5,6}.

Databases

The incorporation of microprocessors into the equipment means that part of the database also has to be in the equipment, otherwise many of the advantages of having the microprocessors there would be lost. The parts that are common to all equipments joined to a process computer should be in that computer, and parts that are particular to an equipment should be in its microprocessor memory.

In the case of LEP, the part that is in the process computer is called the "Equipment Directory" and the part that is in the equipment is called the "Property Module". These two together replace the Data Module of the SPS³.

The equipment directory will be entered by a call from an applications program with name, property, equipment number and possibly value or array. For that name, the directory will give the valid properties, the password requirements, the interface address of the property module, and any translation necessary to call the required property. This directory will contain semi-fixed data, which only needs to be changed if any changes are made to the equipment or to the property modules. The master copy of the data for each equipment directory will be kept in the library, and this master copy will be incorporated automatically into the equipment directory each time it is assembled. The exception to this could be some status indications for the equipments which would be read in from the property modules.

The property module code should be in PROM with only the data table in RAM. This data table will have semi-fixed data, such as conversion factors, tolerances, maximum allowable values and interface addresses, and variable data, such as demanded values,

status, measured values, etc. These distributed data tables form the operational data base for the control system; any program wishing to interact with the equipment or obtain information on its status, will use the property modules and their data tables. There will be no automatic update of a central equipment data base, as with some systems⁷, although copies of the property module data tables will be taken and held in the library periodically. These copies will be solely for archives and for reloading into property modules if, for any reason, the working tables become corrupted.

In addition to this distributed equipment database, ability to record and recall standard settings must be provided. This will be done by programs creating and using files of standard settings kept in the library, which can lead to the automatic check-out and setting-up of LEP for the successive operations of injection, ramping and storage, extending the developments which have already been carried out on other machines^{8,9}.

One of the major advantages of a computer control system is the ability to control derived or virtual parameters, such as Q, beta, etc., rather than individual quadrupole currents. In setting up LEP it will be necessary to have files of a number of this type of parameters from which the settings of individual elements can be calculated, and with which the measured values can be compared. These files will form the central part of the distributed database, and should be partitioned in such a way that a control program will only need to load a subset to provide the environment in which a certain series of actions can be performed.

As the size of a storage ring increases, many parameters become more critical and it becomes more important to provide a good computer modelling or simulation of the beam optics, which takes into account the inevitable imperfections in the machine elements. Considerable computing power will be required for such simulation work, probably beyond that which it would be reasonable to provide in a console computer even in a few years time. One of the advantages of the distributed system that has been adopted is that a big computer, or specialized array processor, can be joined into the network and used by a console computer as an additional resource to which it can send a file, request some operation to be carried out and the results returned. The actual requirements will probably not be known for some time.

Fig. 3 : Database Decomposition

A further requirement for the operation of LEP is to have an installation database, easily accessible from the control consoles, containing the location and interconnection of all the equipment, together with "help" messages telling the operator what to do and who to call if anything goes wrong. This database has to be established early and the information inserted during the construction phase. The relational database system ORACLE¹⁰ has been chosen for this purpose for LEP, and installed on a VAX-780. The proposed hierarchy of databases is shown in Fig. 3.

The Operator Interface

While it is certain that the existing operator interface equipment will be subject to a continuous series of improvements, for example visual displays will have higher definition, full range of colours and be simpler to program and the CRT's may be replaced by flat screens, it is not clear whether any basically new devices will be adopted. Some people prefer the "mouse" to the tracker-ball, but this is only a variant on the same theme. More striking would be the use of the type of eye movement sensors being introduced into aircraft "head-up" displays, which could be used to input to the computer the chosen position on a display, instead of using a cursor or light pen - a blink could then take the place of the interrupt button!

There is a tremendous interest in voice input and output devices in other fields. For accelerator control, voice output to warn the operator of fault conditions is already here, even to providing a choice of language to be used in a multi-lingual environment such as CERN. On the other hand, voice input seems a little dangerous in an accelerator or storage ring context - an inadvertent word might dump the beam!

For LEP, the existing devices used on the SPS consoles would be adequate, but a choice does not have to be made for several years, and progress in this field will be watched, and promising new devices will be tested on the SPS system.

The Timing System

One of the problems of a distributed multi-computer system is to maintain synchronism between cooperating processes. This problem can be largely overcome if an independent distributed timing system is provided. For LEP such a system will be connected to all the microprocessor systems which require synchronizing to better than the fraction of a second that can be provided through the computer network. The timing system will distribute clock pulses and event codes, so that the simultaneous ramping of a series of power supplies can be assured by loading into each a table of values together with an instruction such as "on the nth clock pulse after event Q, load in the first value in the table, followed by successive values every m clock pulses".

Such a timing signal could be transmitted on one of the 2 M-bit/sec channels described below, giving a resolution of 0.5 microsecond, which seems adequate for most purposes. However, there are some special requirements for RF control and beam observation which can only be satisfied by taking the timing for these from the circulating bunches.

Integrated Communications Systems

In any large accelerator, there is a number of services that require cables round or across the ring, and between the equipment buildings, the control centre and the experimental areas. These include inter-computer data links, computer to hardware

connections, audio communications, television pictures, waveform signals, timing signals, access control and interlock signals, etc., which normally employ separate cables for each purpose.

With the distances involved in LEP, it is essential to minimize the number of cables that have to be pulled around the ring, and this can be done by multiplexing as many signals as possible on one cable. We are looking into both frequency division multiplexing and time division multiplexing. The latter seems particularly interesting if we can use the systems being designed to the CCITT recommendations for trunk telephone lines.

In these systems, the analogue signal from the telephone is sampled at a sufficiently high rate (8 kHz) to be able to reconstruct the highest frequency normally transmitted (about 3.5 kHz) and the amplitude of each sample digitized to an 8-bit level, giving a 64 k-bit/sec result. Thirty-two of these signals can be interleaved to give a 2048 k-bit/sec signal, and four of these combined in an 8.5 M-bit/sec channel. Further concentrations of four channels into one give rates of 34 M-bits/sec and 140 M-bits/sec which are already in operational use and 580 M-bits/sec links are being experimented with, using both coaxial lines and optical fibres for transmission.

Such a system will be used to provide 2 M-bit/sec data links for the computer network, the timing signals, a number of 64 k-bit/sec channels for various purposes, several voice channels, and possibly digitized analogue wave forms or digitized slow-scan TV on the same cable. The links between the experiments and the computer centre will also be integrated into the system, possibly using an 8.5 M-bit/sec channel for data transfer.

Most of the links will have to be in coaxial cable, as they will have to go round the ring, where the level of synchrotron radiation precludes the use of optical transmission. However, some 60 kV mains supply cables have to be laid in trenches from the main SPS electrical sub-station to three of the access points on the ring, shown as the lines going across the ring in Fig. 2, and the opportunity will be taken of laying some optical fibre cables at the same time.

Software

The main applications programming language for LEP will be NODAL¹¹. Since the SPS was built, sufficient accelerator laboratories have adopted the interpretive approach for it to be considered conventional these days. The ability to include functions written in other languages and compiled largely overcomes the main disadvantage of an interpreter, the slow speed. It is expected that many of the more complicated functions, such as used in simulation and modelling programs, will be written in FORTRAN.

It has been decided that system software for the multi-microprocessor assemblies will be written in MODULA 2, with a minimum of assembler. This language, together with PASCAL, compiled NODAL and assembler will be used for programming the microprocessors in the equipment. CERN is standardizing on the M 68000 series for 16-bit applications and the M 6809 for the 8-bit systems.

Reliability

With so many components in a control system for a machine the size of LEP, their reliability is of paramount importance. It has been shown³ that the down-time of a large machine like the SPS, due to

failures in the control system, can be kept below 1% of the scheduled time, without appreciable redundancy in the system. However, the momentary failure that results in the loss of one cycle in a pulsed machine may require the refilling of a storage ring, and so the requirements for reliability are more severe.

Reliability can be improved in a number of ways. Conservative designs, quality control, extensive testing and "burn-in" can all affect the reliability of an individual unit. Redundancy, secure electrical supply and "graceful degradation" can improve the reliability of a sub-system. Availability can be improved by built-in surveillance and test programs to speed up the diagnosis of faults or incipient faults, and by modular design to speed up the replacement of faulty items.

Redundancy involves extra cost, and so a careful analysis has to be made to see where it can be most effective. In exploring the other roads to reliability, some conclusions have already been reached. One of the least reliable items in a crate of electronics is the cooling fan, and so the equipment will be designed to use convection cooling wherever possible. Many integrated circuits are becoming available using variations of the CMOS process, giving much lower power consumption with little loss of speed, and this not only helps with the cooling but also allows us to consider the use of individual sealed batteries, float-charged from the mains, for the rack power supplies.

The remote operation and surveillance of the services, such as electricity, cooling, ventilation, etc., is essential when distances up to 10 km are involved, and the control system must continue to provide the required facilities even when the machine is shut down. To cope with this requirement, the LEP network will have a separate node, with links to the process computers controlling these services, to provide a secure sub-net which can be kept operational even if other parts of the control system are shut down as the result of modifications or additions to the machine.

Timescale and Conclusion

A machine like LEP cannot be built quickly, and the boring of the tunnel and completion of the underground works will take four years, and even though installation can start in some areas before all the civil engineering work is completed, it will still be between five and six years from breaking ground to getting a beam round the tunnel. This sets a problem as when to freeze the details of the control system. What today needs a complete printed circuit card may be available tomorrow as one or two chips, and so one would like to leave it as long as possible before freezing the design. On the other hand, it cannot be left too late as it takes some time to have the thousands of units manufactured, and then installed and commissioned with limited staff. Our policy is to have a full scale prototype modular system working in about 18 months time, using present day technology, and then to take advantage of the possibilities of improving or simplifying individual modules before the bulk orders have to be sent out.

As can be seen from above, some principles have been laid down and some decisions made, particularly in the areas in which long-term manpower commitments are involved, and it is aimed to provide a flexible system which can adapt as the circumstances change.

References

- 1 R. Keyser, J-P. Koutchouk, C.R.C.B. Parker, J. Pett and P. Wolstenholme, 'Plans for the LEP Control System', IEEE Trans. NS 26, 3182 (1979)
- 2 G. Plass, 'The LEP Project - Status and Plans', this conference invited paper B3
- 3 M.C. Crowley-Milling, 'Experience with the Control System of the SPS', CERN 78-09 (1978)
- 4 J. Altaber and R. Rausch, 'Truly Distributed Control, Using One Microprocessor per Real-time Task', IEE London, Pub. No. 298 (1982)
- 5 J. Altaber, M.C. Crowley-Milling, P.G. Innocenti and R. Rausch, 'Replacing Mini-computers by Multi-microprocessors for the LEP Control System', this conference, paper W34
- 6 J. Altaber and R. Rausch, 'A Multiprocessor Bus Architecture for the LEP Control System', this conference paper W32
- 7 R. Melen, 'The PEP Instrumentation and Control System', Proc. 11th Int. Conf. on High Energy Accelerators, Birkhauser, 408 (1980)
- 8 J. Gamble, J-Y. Hemery, D. Kemp, R. Keyser, J-P. Koutchouk, P. Martucci, L. Tausch and L. Vos, 'Towards Full Automation of Accelerators through Computer Control', Proc. 11th Int. Conf. on High Energy Accelerators, Birkhauser, 427 (1980)
- 9 G. Shering and V. Hatton, 'Operational Experience using the SPS Control System', IEEE Trans: NS 28, 2225 (1981)
- 10 E.M. Dieckmann, 'Three Relational DBMS, Damation, 137, Sept. 1981
- 11 M.C. Crowley-Milling and G. Shering, 'The NODAL System for the SPS', CERN 78-07 (1978)