

SLAC - PUB - 3624 REV.
June 1985
(A)

CONTROL PROBLEMS IN VERY LARGE ACCELERATORS*

M. C. CROWLEY-MILLING
Stanford Linear Accelerator Center
Stanford University, Stanford, California 94305

Introduction

There is no fundamental difference of kind in the control requirements between a small and a large accelerator since they are built of the same types of components, which individually have similar control inputs and outputs. The main difference is one of scale; the large machine has many more components of each type, and the distances involved are much greater. Both of these factors must be taken into account in determining the optimum way of carrying out the control functions. Small machines should use standard equipment and software for control as much as possible, as special developments for small quantities cannot normally be justified if all costs are taken into account. On the other hand, the very great number of devices needed for a large machine means that, if special developments can result in simplification, they may make possible an appreciable reduction in the control equipment costs.

It is the purpose of this report to look at the special control problems of large accelerators, which I shall arbitrarily define as those with a length or circumference in excess of 10 km, and point out where special developments, or the adoption of developments from outside the accelerator control field, can be of assistance in minimizing the cost of the control system. Most of the first part of this report was presented as a paper to the 1985 Particle Accelerator Conference. It has now been extended to include a discussion on the special case of the controls for the SSC.

The majority of the problems introduced by the size of large accelerators can be divided into two types, those coming from the physical dimensions of the machine, and those coming from the very large number of components to be connected up and controlled.

Extended Version of Paper Presented at the 1985 Particle Accelerator Conference,
Vancouver, B.C., Canada, May 13-16, 1985.

* Work supported by the Department of Energy, contract DE-AC03-76SF00515.

Signal Cabling

One of the obvious dimensional problems is that the lengths of many of the cables, particularly those that have to go round the machine tunnel, increase proportionately and, if nothing is done to reduce their number, the cabling costs for a big machine can become very large. The magnitude of this can be seen from one example. The cost of the signal cables installed initially for the SPS at CERN was about 6 million Swiss francs. If the same system were to be used for LEP, the cable cost would be of the order of 30 MSF (about \$15M at the time this estimate was made) due both to the increased cable lengths and the increased number of components.

The main means of reducing the number of cables is by multiplexing. In many existing control systems, signals from a number of devices belonging to the same service are multiplexed onto a single cable. This can be extended into multiplexing signals from devices belonging to different services onto a single cable. The services requiring signalling cabling round the machine can be divided into two types, those mainly analogic, such as audio and video communications and waveform signals, and those mainly digital, such as links between computers and equipment, terminals and other computers, as well as status and interlock information. There are a number of methods of allowing different services to share the same cable but these usually either require the analog data to be digitized, or the digital data be used to modulate an analog signal. These methods can be divided into two main classes, designated by the method used to share the cable. These are known as Frequency Division Multiplexing (FDM) and Time Division Multiplexing (TDM).

The main difference between these two is that FDM gives each service the full-time use of a portion of the total bandwidth of the transmission medium, while TDM gives each service the full bandwidth for a portion of the total time. FDM has been used mainly for the transmission of analog signals which are used to modulate an RF carrier, to produce an "upconverted" version of the analog signal, occupying a band of frequencies adjacent to the carrier frequency. By modulating each analog signal with a different carrier frequency, and using bandpass filters in the receiver/demodulators, the signals can exist together on a common cable without interference. Many of the present long distance telephone trunk lines use this method of compressing a number of telephone circuits on a single cable. Another form of FDM is the cable television (CATV) system, which carries a number of different programs, each occupying a 6MHz slice of the total cable bandwidth, normally 300 MHz. Digital signals can be transmitted by an FDM system, usually by frequency or phase modulating a carrier.

In contrast, TDM is used mainly for the transmission of digital signals. The time-sharing of the available bandwidth can be carried out in a number of different ways; some synchronous, such as fixed time-slots allocated to each service, and others asynchronous, such as the carrier-sense multiple access/collision detection (CSMA/CD) or token-passing systems used in the growing number of Local Area Networks (LANs) for interconnecting computers, etc. One form of TDM that is of particular interest for a large accelerator is the system being used for almost all new trunk telephone lines. In this, the analog signal from the telephone is sampled at 8 kHz, each sample being digitized to an 8-bit precision, giving a 64kHz digital signal. For the next stage of multiplexing, the Public Switched Telephone Network (PSTN) standards differ between the USA and Europe. The North American standard (AT&T T1) interleaves 24 of these signals into a 1.554 Mbit/s channel, while the International standard (CCITT) interleaves 32 signals into a 2.048 Mbit/s channel. Further multiplexing in both systems is carried out by interleaving the bits from a number of channels, adding some control bits and producing a signal at a correspondingly higher bit-rate. This process is repeated until the full bandwidth of the transmission medium is utilised. In the case of the CCITT standards, four channels are multiplexed into one at each level, but the number varies at the different levels for the T1 standard, being x4, x7 & x6 for the levels so far defined. The resulting bit rates for the two systems are given in Table I.

Table I

MUX Order	Rates - Mbits/s		Cable Type
	CCITT	T1	
1	2.048	1.554	co-ax
2	8.448	6.312	co-ax
3	34.368	44.736	co-ax or optical
4	139.264	274.176	co-ax or optical
5	564.986	—	optical

Optical fibre cables are being used for new installations where heavy traffic is expected, and distances of up to 50 km between repeaters have been achieved at T40 Mbit/s. It has been reported that, since devolution of the telephone services in North America, an increasing number of US companies are building equipment to the CCITT standards.

The use of this type of TDM system for an accelerator control system is of particular interest for a big circular machine where the main equipment to be interconnected is situated at relatively few, well spaced, areas round the machine. This is because the signals going round the ring have to be demultiplexed and remultiplexed at each tap-off point, and this becomes uneconomic if it has to be done too frequently. Connections can be made at any of the multiplex levels, so that 64 kbit/s channels can be used for telephones, terminals, etc. and 2 or 8 Mbit/s for requirements needing a greater bandwidth, such as intercomputer networks. For such applications, the distribution of data signals from a tap-off point to equipment over the area served by that tap can be done by means of one of the standard LANs. For example, in the case of LEP, where this type of TDM is to be used,^[1] it has been demonstrated that a token-passing ring LAN of the type being developed by IBM can be extended over many kilometres by using an 8.5 Mbit/s channel to carry the 4 Mbit/s Manchester encoded data. Standard multiplexor/demultiplexor (MUX) equipment from different telephone equipment supply firms were used for this test, which emphasises one of the advantages of this system; highly reliable and well tested equipment can be bought from a number of manufacturers, and interconnected freely, due to the strict standards imposed.

In the case of a linear machine, where the equipment is distributed fairly uniformly along the accelerator, the FDM type of system may be more appropriate, since taps and relatively simple modulators/demodulators (modems) can be placed wherever required along the cable. The disadvantage of this type is that, since it is an analog signal that is transmitted, high quality cable has to be used, and low distortion, wide-band amplifiers have to be inserted at relatively short intervals. Optical cables are not used so far for this type of system, due mainly to the difficulties of making simple optical taps and wide-band linear transducers. A mixed system is installed for the SLC project at SLAC, using standard CATV cables and repeaters, but with some of the channels used for TDM data.^[2]

The use of optical cables, although attractive from the point of view of high bandwidth and low attenuation, has some limitations due to their sensitivity to radiation. Quite a low dose causes discolouration and an increase in the attenuation. Work is going on to increase the radiation resistance of optical fibres, but so far the dose they can stand is far below that for normal copper cables. The situation is particularly bad for large lepton storage rings, where the synchrotron radiation level is sufficiently high to preclude their use in the machine tunnel. Such use of optical cables may be possible with hadron storage rings, where synchrotron radiation is negligible and beam losses have to be kept low. An alternative for large circular machines is to use microwave links between

the main concentrations of equipment, to carry TV signals and TDM data, but the variation of propagation delay with weather conditions and the possibility of interference must be taken into account, particularly if the links are to be used to carry precise timing information.

Distributed Computing

Another result of the dimensions of a large accelerator is that the computer system for control should be distributed around or along the machine. It should be a fully distributed system in all aspects, with appropriate devolution of activities. Some systems in use for existing accelerators that have distributed processors use these only as slaves for a central computer, which is involved in all main control processes, including the use of terminals for the local control of equipment remote from the central system. With a very large accelerator, even with devolution of some activities to other processors, a single central computer can become a bottle-neck which limits the number of simultaneous real-time tasks that can be carried out. This limitation is usually imposed by the standard operating system or by the I/O and interrupt capabilities, rather than by raw CPU power. The central control facilities should be supplied by a network of computers, some of them being dedicated to specialized tasks, such as running consoles, carrying out alarm analysis or acting as a file-server, while others perform different tasks as required for the particular operation being performed. The high-speed network connecting these computers should preferably be extended to the distributed computers controlling those local processes that need rapid access, such as the beam instrumentation, giving direct communication rather than through a gateway between two separate networks.

In a fully distributed system, with remote computers able to carry out autonomous tasks and provide full support for local control of both local and remote equipment, any bottle-neck can be overcome by adding another computer to the network to take over some of the tasks, without disturbance to the rest of the system. The value of this has been shown at the SPS at CERN, where the conversion from a pulsed proton accelerator firstly to a proton-antiproton storage ring and now to an electron-positron injector for LEP, with the considerably increased control requirements, has been achieved by adding computers, with their associated data links and interface equipment, without disturbing the existing system, because the distributed operating system allows direct high-speed communication between any two computers. The same type of system will be used for LEP.^[1]

Interface System

Not only does a large accelerator have many more of each type of component than a small one, but additional systems have to be connected to the control system. For example, with a small machine such things as the power distribution and water-cooling systems are usually manually controlled, and the accelerator is totally within a guarded site. With a large machine, virtually everything has to be connected to the control system, including the surveillance of access to the distributed plant buildings, unless the project can afford to have round-the-clock personnel at every site.

The method chosen for connecting up the many thousands of components to the control computers has a great influence on the cost of the control system. As mentioned earlier, a small machine should use standards wherever possible, and the most popular standard interface system in the accelerator field is CAMAC. Although CAMAC crates and modules are relatively expensive, and usually require junction and distribution boxes to match the standard modules to the monitoring and control requirements of the individual pieces of equipment, it is at present the most economical solution for a small machine, if development costs are taken into account. The situation changes as the size of the machine increases, and the design of special modules to reduce the interface costs becomes advantageous. This has been done for some medium sized machines, resulting in systems such as the MPX interface for the SPS and the PADAC interface for PETRA and HERA.

All these interface systems were developed before the advent of cheap, powerful microprocessors. In today's situation, and especially for the large machines, we can take a quite different approach to the interface problem. To illustrate this, let us take the case of the control of a large power supply. The older type of supply would have some relay logic in it, for interlocking and possible sequencing of switch-on operations, an analog feed-back system for stabilizing the voltage or current and possibly a hardware function generator for ramping. It would be controlled and monitored by a number of separate digital and analog input and output circuits from the interface system. Now a microprocessor system can be incorporated into the power supply unit at the design stage, and can be programmed to perform all these duties. In addition, it can run periodic test programs to check the health of the power supply and produce diagnostic information in the case of failure. Only a single connection is now needed to the control system, since the microprocessor can send and receive messages replacing the information carried on the separate circuits previously. In the case of a ramped power supply, a connection to the timing system is also needed, but this can be provided in the same cable. The control connection to the equipment

can be an interface to a data highway which goes round the equipment to be controlled from the local computer. This data highway can be relatively simple and cheap, and does not have to have all the facilities of the LANs intended for computer interconnection. Suitable systems for this application are presently available, such as MIL-STD 1553 and Intel's BITBUS, and similar systems, with single chip interfaces, are expected in the near future.

An interface system that requires some processing power in the equipment can certainly be justified for a component with the control requirements of a power supply, but there are many components with very simple requirements, where it might be thought that the incorporation of a processor could not be justified. However, single-chip microcomputers are already available with on-chip I/O devices, such as ADCs and serial links, together with a small amount of RAM and firmware that includes an interpreter. Where the quantities are large enough, it seems to be economically advantageous to have a single chip designed to interface directly to the chosen data highway, and to provide the control requirements for a simple piece of equipment, or for a number of independent status or control actions grouped together. One chip could cover the control requirements for a number of different types of equipment, with different programming. These programs could be burnt into "piggy-back" PROMs that plug into the chip carrier. The power for the chip could be carried on the data link cable, especially if the CMOS process is used, as the power requirement would then be small, and the chip could be mounted inside the plug on the end of the highway cable, as shown in Fig. 1. This would then form an "intellegent" plug which could either control simple devices directly, or exchange messages with more complicated devices incorporating microprocessors. One could then have a uniform message-orientated interface system, with every piece of equipment being able to test itself, or be checked by the interface plug, and the result reported to the controlling computer. The importance of this will become evident below, in the discussion on reliability.

Modeling and Simulation

Another effect of the quantity of components in a large machine is the increased requirements for real-time modeling or simulation of the optics for control purposes. The larger the machine, the more critical its adjustment becomes, as more modes of instabilities or resonances become possible, and the smaller the deviation from the ideal orbit that can be allowed. Thus computer modeling, which provides operational convenience for small machines, becomes essential in large ones, to determine the adjustments necessary to achieve an adequate dynamic aperture: random adjustments can too easily result in a divergent situation.

However, unless new techniques are developed, the complexity of the computations required for modeling goes up rapidly with the number of components that can affect the beam, with a corresponding increase in the computer power needed to carry them out in a time short enough to be useful in setting up the machine. The distributed control system should be provided with sufficient power to gather and format the beam diagnostic data in less than a second, but it may then be necessary to pass this data to a specialized computer or array processor, which is either part of the control system or a facility of the central computing services of the laboratory. In a recent estimate of the computing needs for the SSC,^[3] the need for the full time use of a computer of the class of a CRAY 1S for the design calculations was expressed. If such a machine is obtained for this purpose then, once the design is complete, it could be coupled to the control system network for the on-line modelling activities.

The criticality of the parameters of a large machine also means that, in the case of a ring where the injected beam has to be accelerated up to the storage energy, the tolerances on the tracking of the various elements during ramping have to be very tight, and this must be taken into account in the design of the power supplies, interface and timing systems.

Operator Interface

The problems of providing an operator-friendly interface for the control operations are not fundamentally different for a small or large accelerator, and the actual methods and devices used for interaction are equally open to personal choice, but the greater distances mean that more facilities for control local to the equipment distributed over a large area must be provided for commissioning, test and maintenance. Obviously, local consoles must be provided for the larger concentrations of equipment, but there is also a need for small portable control units that can be carried round a ring, or along a linac, for tests and checking. It has already been shown that "leaky" cables can be used for radio communication between hand-held stations in a tunnel and outside, and similar means could be used to link a portable terminal to the control system for checking equipment in the tunnel, without having to provide plug-in points at frequent intervals.

A problem with any widely dispersed system that provides the possibility of taking control of part or all of the equipment from many different locations is how to keep track of what people are doing, particularly those outside the main control room, and how to prevent unauthorized access to the system. Many schemes for passwords, locks and keys, etc., have been proposed and some implemented, but no fully satisfactory solution has yet been demonstrated, and work is needed on this subject.

Data Bases

There are a number of data bases involved in the construction and operation of a large accelerator, and nearly all of them have to be used, directly or indirectly, by the control system. The first data base to be formed is that for the optical design of the machine. Many design and ray tracing programs have been written in the past, with many different formats for the data files. A proposal has been made for a standardised format,^[4] and this is being adopted by many laboratories. Extracts from this data base are required by the control system for the on-line modelling programs.

Next in chronological order is the machine constructional data base. For a large accelerator, it will be essential to have a data base in which is recorded information on all components of the machine, their parameters, procurement and installation schedules, interconnections and cable routes, etc. Such a data base will require a commercial data base management system, preferably of the relational type, running on a large computer with access from a very large number of terminals. This will need to be closely linked with the CAD/CAE systems

used for producing the drawings for the accelerator, to ease the task of cross-correlation. The control system will need to have access to this large data base, both during construction and in the operational phase, but access to the data is likely to be too slow for any of the data to be used directly by on-line programs. For this purpose, it will be necessary to extract the relevant information and store it in a suitable format in the central control data base.

Another data base that should be accessible to the operators of the accelerator, but not directly involved in control, is a subset of the personnel records, giving the persons responsible for the various parts of the machine, together with addresses and telephone numbers, etc.

The control system data base should be distributed amongst the various computers in such a way as to minimise the network traffic in performing the control and surveillance activities. The central data base should contain all the non-transient data for the whole machine, for down-loading or backup purposes, but the working data base should be distributed round the system. This not only allows subsystems to be operated autonomously, but also avoids the heavy traffic that would be involved in updating a single central database. If all the computers, including the central ones, are on the same network level, allowing high speed communication between any two computers, any process that requires data can obtain it from the computer keeping that part of the data base.

A central data base is required for data concerning the machine as a whole, such as the optical model, standard configurations and timing information. This could be held in a central file-server or in the computer used for modelling and simulation. Such a central database should also hold the excerpts from the relational data base mentioned above.

At the lowest level, each of the microprocessors in the equipment or interface must have a small data base of I/O addresses, conversion and calibration factors, etc., if the communication with them is to be in the form of messages, as proposed in the section on interfaces.

Timing

It is now common practice to distribute round an accelerator a timing signal consisting of clock pulses interspersed with coded information on events occurring in the machine cycle, such as injection, start of ramp, beam dumped, etc. The timing system producing this signal train should output some of the events at relatively constant times for each cycle, while provision must be made for others, such as that signifying that the beam has been dumped, to be inserted at any time they might occur. Equipment can be programmed to utilise the signal train to perform some operation at a certain number of clock pulses after a given event.

The additional problems of timing in a large machine are connected with the relatively long time for a particle to pass along or round the machine. For single pass machines, such as colliding linacs, the problems do not significantly increase with size, as "time" effectively travels down the machine with the beam. The absolute precision of launching the beams into the linacs to make them collide at a given point is unchanged, but the greater distance that the timing signals have to be transmitted, the more difficult it is to prevent the jitter becoming excessive.

In the case of a large particle/antiparticle storage ring, it is necessary to consider "time" either as something that travels round the ring in opposite directions for the two beams, or as a series of events occurring simultaneously at all the intersection points.

Programming Languages

Although the choice of programming languages does not seem at first sight to be a matter dependant on the size of the machine, it can be an important factor in efficient operation of the fully distributed type of computer system that is recommended for a large machine. In such a system, a strong case can be made for the use of a control language that can be interpreted for the applications programs. An interpreter is an on-line linker, which takes the source code directly and, according to the commands, searches for and links together code segments usually known as functions. In its simplest form, as in interpreted BASIC, the functions are limited in number, and relatively small in size, performing simple actions such as addition, multiplication, looping, etc. A program written in this type of BASIC consists of a number of commands, each of which has to be found and linked each time the program is run, and so the execution is relatively slow. On the other hand, if an error occurs in running such a program, the statement producing the error can be displayed without difficulty, since the source code is interpreted directly. Thus debugging is made much easier.

At the other extreme, an interpreter can be written for a control language in such a way that, as well as giving the usual facilities, it can be used to link an arbitrarily large number of functions that are, in principle, unlimited in size. These functions, or sub-programs, can be written in any suitable language and compiled, so that the speed penalty of having them linked on-line can be quite small, but the flexibility this on-line linking gives in a fully distributed computer system is very valuable. In addition, means can be provided for the interpreter to control the import and export of program segments, and the synchronization of programs running in different computers. Programs can then be stored centrally and loaded dynamically as required, not just at initialization time, and parts of a program can be exported to other computers for execution as required. Such an interpreter was written for the control of the SPS, using the accelerator control language NODAL. Simple programs can be written in NODAL and executed directly. After debugging, parts of the program can be compiled if the speed of execution is insufficient. For more complicated control programs, the problem is usually divided into three parts, data acquisition carried out by the interpreter, manipulation of the data by a function or sub-program written in such a language as FORTRAN, and then output of the results by the interpreter using display or setting functions. NODAL, or one of its derivatives, has subsequently been used at DESY, JET, TRISTAN and the Rutherford Laboratory.

One of the disadvantages of interpreting the source code directly is that comments and documentation written in the program increase the size of the files that have to be loaded into the computers that are to execute them and, when there are limitations on buffer size or data link speed, this tends to encourage minimum comments and poor documentation. This can be overcome by insisting on full comments and in-program documentation, and providing a source code management system which, as well as storing the complete programs as written, also takes them and strips the comments, abbreviates where possible and stores the resultant code in an execution library. A call for listing would display or print the original with the full documentation. If this should be edited, the result could be stored as a new source module, but a stripped version of it would not automatically replace the execution module until some positive action were to be taken, and this might be under the control of the operations group in a running machine.

Errors and Alarms

The basic problems of error and alarm indication are not dependant on the size of machine, but the greater number of components in a large accelerator make it more important to solve the problems described below, and avoid the operators being overwhelmed by alarm messages (The Three Mile Island effect!).

In a fully distributed system, it is essential to have a single computer designated as the recipient of all alarm and error messages. The different types of messages can be divided into several levels. At the highest level are the primary faults causing stoppage or reduced performance of the accelerator, which must be identified and brought to the notice of the operators in a clear and unambiguous manner. At the next level are the consequential faults, where equipment cannot operate correctly because of the primary fault that has occurred in some other device. A simple example of this is when a string of power supplies trips off because of failure of the cooling water supply to the magnets they are supplying. The cooling water supply is the primary cause, and should be indicated as such. However, when this is restored, the tripped-off power supplies become primary faults, unless they have provision for automatic switch-on when the water flow is restored.

The problem of sorting out the consequential from the primary errors, to avoid the situation where the operators are given so many unnecessary alarm signals that they ignore the important ones, is one that has not been completely solved either in accelerator control systems or in the nuclear power industry, despite considerable investment. It may be that some of the recent developments in AI (artificial intelligence) techniques may show the way to solve the problem. A self-learning system was proposed for the SPS many years ago, but the techniques then available were inadequate.

The next lower level of error messages come from the surveillance programs that detect that some parameters are drifting out of the pre-set tolerances. These should be indicated to the operators in a different format, or on a different screen, to the primary alarm messages. Then, at the lowest level, there are the transient failures, such as errors in message transmission, which are detected and corrected by retransmission, without affecting the operation of the accelerator, but which need to be recorded for off-line analysis to detect equipment that may have intermittent faults, or be operating in a marginal fashion.

Reliability

As the size of a machine and the number of components increases, The reliability of the individual components must increase, if the operating efficiency, defined as the ratio of actual to scheduled time of operation, is to be maintained. An average efficiency of 90% can be achieved with existing medium sized accelerators and as long as the control system faults are responsible for less than a tenth of the total down time it is normally considered to be sufficiently reliable.^[6] The reliability of a system is most affected by the reliability of the weakest components, but which those are can often only be found out from experience. Therefore, an attempt should be made to use well-tested equipment where possible, such as the PSTN TDM multiplex system mentioned above, or to build up a test system sufficiently early to get operational experience with new equipment. In general, reliability of a system can be improved by increasing the mean time before failure (MTBF) of the individual components, and/or reducing the mean time to repair (MTR), which includes the time to diagnose the fault and then to repair or replace the faulty component. In the case of a large accelerator, the time taken to get to the locality of the faulty component can be a significant part of the MTR.

The complexity of electronic equipment, especially computers and allied equipment, has increased enormously over the last decade, but this has not been accompanied by a decrease in reliability; on the contrary, reliability has generally increased. There are several reasons for this, one of the main ones being the increased use of large scale integration, allowing a reduction in the number of separate components and their interconnections, particularly those made by means of plugs and sockets, which are frequently the source of trouble. A computer that would have needed a complete crate of separate boards a short while ago can now be fitted on to a single board, and will shortly be available on a single chip. The use of large scale integration has also allowed a certain amount of redundancy to be incorporated to overcome faults, such as in self-correcting memory.

One way of increasing the reliability of a control system is by duplication or triplication of the equipment, but this is normally not economically justifiable for an accelerator system, with the possible exception of those parts involving protection of personnel. As computer hardware nowadays is no less reliable than the interface, cabling and sensors, duplication of the whole system is needed to get appreciable gain. Instead, the system should be examined for possible sources of failure, and means of eliminating or reducing them sought, to increase the MTBF. This can include making sure that components are run well within their ratings and that careful attention is paid to the thermal environment, as

the MTBF of electronic components increases if they are run well below their maximum allowable ambient temperature. Even small doses of radiation can reduce the reliability of integrated circuits, and every attempt should be made to keep electronic equipment out of the radiation areas. Where it cannot be avoided, it may be possible to use radiation-hardened components, but even these become unreliable after doses in the range of ten to a hundred kilorads. The extra expense of using components to military specifications generally in the control system is not likely to be justified, unless the price differential is reduced substantially in the future.

An attack on the MTR can also give significant results. Very often the diagnosis of a fault takes longer than its repair, which emphasises the importance of comprehensive test and diagnosis programs. As mentioned above, the provision of processing power in the equipment and interface allows diagnostic programs to examine the functioning of the system down to the lowest levels. All equipment should be in modular form for easy removal, and recovery from a fault should be by exchange of modules, rather than repair on the spot.

Electronic equipment cannot operate reliably with unreliable power, and this is another area which requires careful attention. With a large accelerator, the control system is required to monitor and control the electrical distribution system, and so at least part of it requires a secure supply, with sufficient autonomy to maintain safety systems and carry out recovery procedures in the case of a major power outage.

As more and more integrated circuits are becoming available in CMOS form, without significant loss of speed as far as control applications are concerned, but with a small fraction of the power consumption of the present equipment, it becomes realistic to consider powering parts of the control system by sealed rechargeable batteries. Normally, the batteries would be floating, the power being supplied by a charger from the mains. The size of the batteries could be chosen according to the time of autonomy required for different parts of the system. Sealed batteries are already available with 5 to 10 years guaranteed float life and they can withstand up to 2000 discharge/charge cycles, depending on the depth of discharge. An important additional advantage is that the use of floating batteries for supply also increases the reliability of the electronic systems, since they can reduce both voltage variations and mains-borne noise to a negligible level.

Experimental Data

With a very large machine, the data links needed for the experiments should be an integral part of the general machine communication system and not something added on later, as has so often been the case with existing machines. There are two main services required by an experiment remote from the main computer centre. The first of these is for terminal access to the central machines to run programs and get the results back as tables or displays. This service does not require very high-speed links. The second has traditionally been for transmitting sample data to the computer centre for analysis to check on the operation of the experiment. There seems to be a tendency to provide enough computer power at the experiment to do this locally, but if this is not the case, it will be necessary to provide suitable links. The data rate required is likely to be high, even though a considerable amount of reduction may be carried out in the experimental data acquisition system. If the TDM system described above is used, there should be no difficulty in providing an 8 Mbit/s or 32 Mbit/s channel to the computer centre for each experiment.

A Practical Case

Some of the problems particular to a large accelerator have been discussed and a number of solutions proposed. None of these call for any radical new developments, but assume a continuing progress in integrated circuits and similar devices. To proceed further, and make choices between alternative possibilities, particular cases have to be considered. There are two machines that are being studied at the moment, a 1 Tev on 1 Tev linear collider for leptons and a 20 Tev on 20 Tev storage collider for hadrons, the SSC. The study for the linear machine is in a very early stage, but the control problems are likely to be similar to those for the SLC, and the methods of solving them may also be similar. The study for the SSC is much further advanced. A report has been published,^[6] and it is hoped to make a decision next year on the type of magnet to be used, which will allow some of the major parameters to be settled. Let us look at the various subjects that have been discussed above, in this context.

Cabling for the SSC

The choice of the type of multiplexing to be used for the SSC will depend on a number of factors, one of which is the overall cost. Sufficient data is not yet available to make detailed cost estimates for the two types, even if only one of the three reference designs is considered, so it will not be attempted, but a few non-cost factors that may influence the choice will be discussed.

The broadband CATV FDM system is in sufficiently widespread use that many reliable and cheap components are available for one-way multiplexing of analog signals, particularly if they fit in the standard TV channel allocations. Duplex operation can be obtained either by using two cables, or by using half the available bandwidth for each direction, with frequency conversion equipment at the "head" end of the cable. Modems are also available for transmitting digital signals, usually by frequency-shift keying of a carrier (FSK). Modems are available at present for rates up to a maximum of 4 Mbits/s, using a 6 MHz channel, but faster ones should be available in the near future, using more than one channel, such as 10 Mbits/s in a 12 MHz band. Development of industrial equipment for the broadband system should receive a strong impetus from the decision by General Motors to impose its use on all its suppliers of equipment for factory automation.^[7]

On the other hand, there is growing interest in the Integrated Service Digital Network (ISDN) concept, which aims to supply all telephone and terminal services for the office of the future, both local and distant, on a single network. This is being backed by the giants of the telecommunications industry, and uses the type of TDM described in the earlier section of this report. This should result in the commercial availability of the additional pieces of equipment needed to use this type of system for accelerator control well within the time scale of the SSC. CERN is having to develop some of this equipment to use the TDM system at LEP, because of the shorter timescale.

Using the TDM system, analog signals have to be digitized. In the case of many existing accelerator control systems, waveform signals are already digitized for transmission over any appreciable distance, so that is not a problem. The main difficulty arises with TV signals. A digitally encoded TV raster, with the normal frame repetition frequency of 60 Hz, requires a 60 to 80 Mbit/s channel, and even with this, some resolution is sacrificed. However, slow-scan and compression schemes have been developed that can transmit acceptable pictures for many of the purposes for which TV is required round an accelerator, such as for surveying personnel entry doors, using a single 2.048 Mbit/s TDM channel.

A point to note is that, in the case of many existing machines, the

responsibility for the control system, the personnel protection system, the telephone system and the "emergency off" system has been divided between different groups of the project team, and that has resulted in separate cables being laid where a single one would have been possible if all these services had been coordinated by one group.

Interface for the SSC

One of the problems that results from the incorporation of microprocessors in the equipment to be controlled is that the design of the hardware and software for these microprocessor systems tends to become the responsibility of the hardware group concerned, and moves out of the hands of the control system group, unless special arrangements are made to supply the services of controls group personnel to each of the groups concerned. In this case it may be possible to achieve a uniform hardware and software approach, but it is probably unrealistic to hope to do so, especially as, on the time scale of the SSC, much of the equipment available commercially will already incorporate microprocessors following the manufacturer's own standards.

The alternative way to avoid chaos in the integration of the systems, and the subsequent maintenance, is to lay down strict rules for the hardware and software interfaces, and insist that the equipment is constructed in modular form, with diagnostic software to identify a faulty module, so that recovery from a failure is by replacement of a module, without having to find the internal fault.

If the proposal is accepted that communication between the control computers and the equipment should be in the form of messages, even for the simplest devices, then each piece of equipment can be considered as a named device that has certain properties, defined in the software in the equipment microprocessor or in the "intellegent" plugs described above, or their equivalent. If these messages are in the form of a string of ASCII characters, then any piece of equipment can be tested off-line by connecting it to a suitable terminal and exchanging messages.

It is likely that a suitable multidrop highway will be available, with single-chip interface, on the time scale for the SSC, and single-chip microcomputers with on-chip I/O capabilities adequate for the "intellegent" plug are available now, but it may be necessary for the project to make some financial contribution to the development of a single chip that combines the two. The other potential applications for such a device might induce a manufacturer to take over part of the development costs, but even if that is not so, with a requirement for about 45,000 of these, it should only add a few dollars per unit to carry the full cost of development.

A Computer Network for the SSC

A proposal for the control computer network for the SSC is given in the reference design document, and a simplified view of the proposal is given in Fig. 2(a). From an organizational point of view, it may seem logical to chose a geographically determined hierachy like this, but it can result in a heavy overhead for relatively simple operations. For example, if a console calls for an acquisition of the closed orbit position, the main computer has to send a message to each of the 12 computers at the next level to get their portions of the data. These in turn ask the 108 computers at the next level, and these ask the 970 at the lowest level to get it from the BPM interface modules. The data returned at each stage has to be concatenated and passed up to the next higher level. With most standard operating systems, the interrupt and message-handling routines impose a considerable software overhead at each stage.

These difficulties can be largely overcome if all computers are at the same hierarchical level, connected by a high-speed message transfer network, as shown in Fig. 2(b), and if they have an operating system designed for real-time distributed control. In such a system, the call for a closed orbit acquisition from a console computer would be in the form of a "function broadcast" which is recognised and acted on by each of the computers that has BPMs connected to it. The data returned from these would be addressed directly to the console computer concerned, which can then concatenate the data and run the appropriate treatment and display programs.

A first guesstimate would result in about 250 computers on the network, with an average of about 200 of the interface plugs per computer. Present hardware for multidrop highways limits the number of drops per highway to less than 100, so an average of two or three highways per computer would be needed, which fits in quite well for the case of a computer roughly in the middle of a batch of distributed equipment it has to control.

It might be thought that, with all 250 computers on one network, even a very high-speed one, some congestion could occur, and that the network might be the limiting item in speed of operation of the system. Experience with existing systems of this type has shown that the network is seldom the bottleneck and that when delays occur, it is usually in the message handling part of the operating system of the computers. The SPS system has some 50 computers joined by a packet-switching network with a line speed of only 750 kbits/s, and yet measurements show that even at maximum activity the network is only loaded to about 20% of its full capacity. Roughly equivalent loading would be expected from having 250 computers on a network with a line speed of 4 Mbits/s, which

is by no means the limit, even today. However, although the network speed may be adequate to deal with the full traffic, a case can be made for splitting it into a number of separate sub-nets for other reasons. For example, there are a number of sub-systems, such as that for the cryogenics, where the main communication needs are between parts of the same sub-system and the traffic to the rest of the control system is quite light. There are also some services, such as the electrical distribution, the cooling and ventilation and the vacuum systems, that need to be kept in operation when the machine is shut down, and it is convenient to have these services on a sub-net, which can be kept running even if other parts of the network are shut down for the inevitable additions and improvements as time goes on. For this reason, the control network for LEP will have two main rings round the machine, one for machine control and the other for the services mentioned above. The service sub-net in this case will be powered from the electrical distribution substation batteries, and some parts may be duplicated, to ensure maximum availability.

It is far too early to make a firm choice for the actual network for the SSC, as there are many competing systems, and it is not yet clear which will become dominant. The IEEE P802 committees are trying to arrive at standards for four different types of networks, as agreement cannot be obtained on the use of a single type for all purposes. There are sound arguments for using the token ring type (P802.5) for the SSC, but the token bus type (P802.4) seems to be getting a lot of support for industrial process control.^[7] With giants of industry like IBM behind the token ring and General Motors behind the token bus, the fight should be interesting. The SSC can sit on the side-lines and choose the most cost-effective when the time comes. If the token ring type of system should be chosen for the SSC network, the sub-nets could be arranged as shown in fig. 3, where there are a number of token rings going round the machine, carried on separate channels on the same TDM cable, and joined together by bridges. This would allow computers on the same sub-net to communicate between themselves directly, and to communicate with any other computer through one or more bridges, thus preserving the single hierarchical level. If it was decided to use PCAs (see below) for the computers, any computer could be connected directly to more than one sub-net, without having to use bridges, just by plugging more than one Data Link Controller into the crate. For a token bus system, the layout could look almost the same, with a break somewhere in each of the rings, although the method of operation differs considerably between the two systems.

Computers for the SSC

The proposals so far do not put any constraint on the types or sizes of computer to be used, except that they must have compatible operating systems that provide facilities for fully distributed, multicomputer control. However, a number of difficulties have been experienced in the past, and methods of avoiding these should be investigated. The main difficulties have been:

- An accelerator normally has a life longer than one particular model of computer and in the case of several existing machines the control system has had to be extensively rebuilt at least partly as a result of the original computers becoming obsolete.
- The computer manufacturers' standard operating systems are getting larger to provide more and more facilities and to hide the hardware structure from the user, such as in the virtual memory systems. This has led to slowing down of response to program changes, which is a serious disadvantage for the type of real-time activities needed in a control system. Once having chosen to use the standard system, one has to introduce each new version that is brought out, even if the new facilities are not needed, or lose the manufacturers support.
- Once a computer manufacturer has been chosen, one is locked in for all future changes or additions. In most cases one has to use the manufacturer's maintenance service, the cost of which is usually based on a percentage of the cost of the equipment, and which can be quite high if one requires rapid response to out-of-hours calls, which is often the case, since most machines are operated round the clock.

Consideration of these disadvantages, and others, led to the proposal that for LEP most of the distributed computers should take the form of crates of microcomputers, which were called Process Control Assemblies (PCAs).^[1] The idea was that, instead of having a single processor with a complicated multi-tasking operating system, there should be a separate processor for each main type of task, such as network interface, interpretation of programs, local database management, supervisory duties, I/O to equipment, etc. In this case, the operating system reduces to a simple scheduler and supervisor. If a suitable standard crate and bus system is chosen (VME in the case of LEP), then modules from any manufacturer can be used to build up the PCA, as long as they follow the bus and inter-module communication standards.

— The various PCAs can be tailored to suit the different applications by using a suitable mix of the standard modules. Examples from the LEP proposal are shown in Table II.

Table II

Process Control Assembly (PCA) of Microprocessor Modules
to Replace Minicomputers in the LEP Control System

Types of Modules:

DL	Data Link driver
BC	Multi-drop Highway Branch Controller
DC	Display Controller
OI	Operator Interface (Knob, Ball, Touch-buttons, etc.)
DD	Disk Driver
MS	Mass Storage (Semiconductor)
GP	General Purpose processor which, according to the software loaded can perform the following tasks:
NI	NODAL interpreter
SU	Supervisor
ED	Equipment Directory
CT	Computational Task

PCA's for different uses can be built up from various combinations of these modules, as shown below.

Type of PCA	Number of each type of module required									
	DL	BC	DC	OI	DD	MS	NI	SU	ED	CT
Equipment control Min. configuration	1	1	-	-	-	-	#	1	1	-
Equipment control Max. configuration with local console	1	*	1	1	-	1	#	1	1	1
Mini-console	1	-	1	1	-	-	1	1	-	-
Main console	1	-	\$	#	-	1	#	1	-	#
Library	1	-	-	-	+	+	1	1	-	-
Stand-alone test	-	1	1	1	1	-	1	1	1	-

- Notes: * According to number of highways
 # According to number of simultaneous processes
 \$ According to number of simultaneous displays
 + According to amount of storage required

One big advantage of this scheme is that the system can be designed so that if one type of task causes an overload, an additional module of the appropriate type can be plugged into the crate, and the task divided amongst the processors. It also gives the facility, mentioned earlier, for connecting a PCA to more than one sub-net and directing messages appropriately and for driving an arbitrary number of lower level highways to connect to the equipment to be controlled.

While the great majority of computers in an SSC control system could be PCAs of this type, there is also need for some large minicomputers with full operating systems giving a comfortable environment for program development, for running commercial data-base systems, etc. These will have to be coupled to the network, to communicate with the PCAs, but they should not be directly involved in the real-time tasks of running the accelerator. In the same way, a "number cruncher" could be added to the network to provide a service for those programs requiring more power than is provided in a console PCA.

Feedback Systems and Response Time

An important consideration in the design of a computer control system is the response time requirements for the different types of servo and feedback systems. These can vary from the sub-microsecond response required for the RF feedback systems, which is expected to be provided by specialized local hardware,^[6] to the multi-second feedback loops for compensation of the effects of slow beam loss during the storage period.

In between these extremes there is likely to be a number of requirements for feedback systems with response times in the region of a few milliseconds. Some of these will be purely local loops, such as the current stabilization in a magnet power supply, that can be handled within one microprocessor. Other may involve transmission of data over an appreciable distance, and may include more than one processor in the loop. Response times of a few milliseconds can be obtained for communication between two computers if the network joining them has provision for messages of different priorities, and the highest priority is reserved for such messages. Should it be necessary to have a dedicated fast long-distance link as part of a feedback system, it could be provided by allocating a channel in either of the two types of multiplexing schemes described earlier.

Conclusion

In looking at the particular problems arising in the design of a control system for the SSC, a number of the options have been discussed, some proposals have been made and subjects pointed out where additional investigation or development is needed. One of the greatest difficulties in coming to any firm conclusions in the near future is the long time-scale of the project at a time when advances in electronics are proceeding rapidly, and what seems to be a pipe-dream now may well be routine in a very few years time.

REFERENCES

1. "LEP Design Report" CERN-LEP/84-01 June 1984
2. W. Struven "Wide-band Cable Systems at SLAC" IEEE Trans. Nuc. Sci. NS-30, 4, August 1983.
3. R. Diebold "Computing needs of the Superconducting Super Collider" IEEE Trans. Nuc. Sci. NS-32. No.1 Feb 1985.
4. F.C. Iselin "The MAD program" CERN-LEP/TH/ 83-30.
5. M.C. Crowley-Milling "Experience with the Control System of the SPS" CERN 78-09, December 21, 1978
6. "Superconducting Super Collider" Reference Design Study for U.S. Department of Energy May 8, 1984.
7. "Automaker's MAP leads the way to factory automation" COMPUTER DESIGN, May 1985. p53.

FIGURE CAPTIONS

1. "Intelligent" plug for interface to equipment.
2. (a) Computer layout proposed in SSC reference design study.
(b) Layout with all computers at the same hierarchical level.
3. Multiple ring LANs, using common TDM link for long-distance connections.


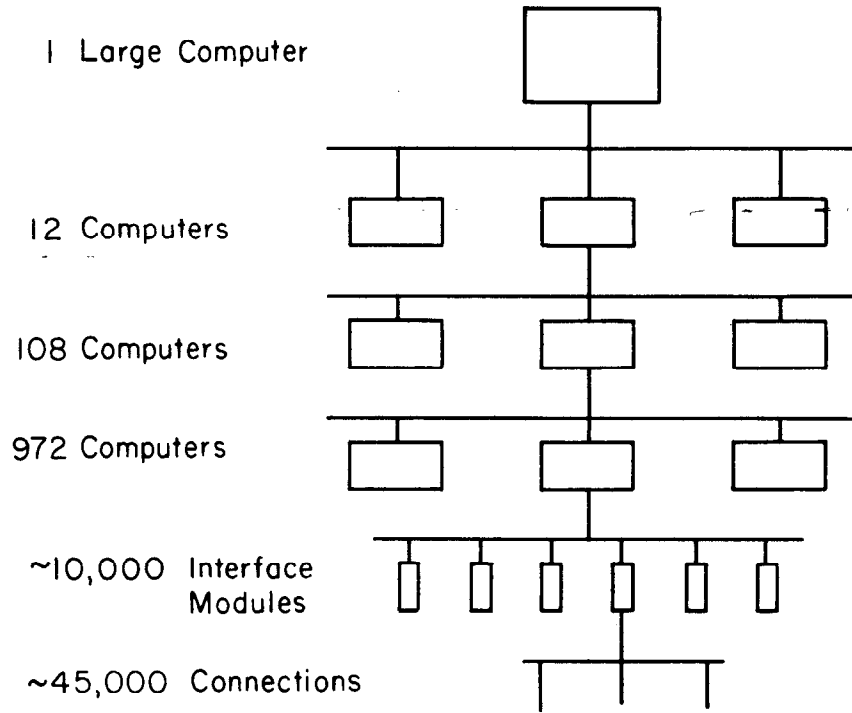



Fig. 1

COMPUTER SYSTEMS FOR THE SSC



(a)



(b)

6-85

5174A2

Fig. 2




Fig. 3