COMPUTER CONTROLS AT THE SUPER ACCELERATOR

Dixon Bogert
Fermi National Accelerator Laboratory

Since the SSC controls may have many similarities to those of the Fermilab Doubler, it is useful to review the Doubler controls. On the other hand, there are some significant differences that are worth a little bit of attention and there may even be some areas that will attract some interested industrial participation.

Although the circumference of the tunnel obviously scales linearly for an accelerator that is twenty times more energetic but using the same magnet technology as the Doubler, there are some aspects of the control system which it is now believed will not scale linearly. In fact, there are some instances where fewer devices will be needed, and others where the number of devices will be relatively the same as for the Doubler. For example, for the SSC refrigeration, the plan is to have only twelve very large distributed refrigerator systems rather than scaling the design of the relatively smaller Doubler refrigerator system. The Doubler refrigerator system already includes at least twenty-four distributed refrigerator engines in individual buildings plus about eight compressor buildings leading to a net count that is already in the thirties. Therefore, this is an example of an SSC system projected to have fewer distributed components than presently in the Doubler. An example of a system whose components will remain rather similar in number as in the Doubler is the correction magnet system. There are "cells" in both the Main Ring and the Doubler that consist of four dipoles, a focusing quadrupole, and then a correction coil package. For the SSC, the number of "cells" as defined by the number of dipole correction elements does not scale by a factor of twenty, but is similar to the current Doubler number. Of course, the number of main dipoles per cell will increase, unless very long dipoles are built. Therefore, the number of correction function generators required will remain at the level of a few hundred, rather than maybe five thousand.

One of the necessary features of the superconducting accelerator that exists at Fermilab is an emphasis on distributed processing. The need for distributed processing was driven by considerations which were only beginning to become important for an accelerator of the physical dimensions of the Doubler. These were especially important for the quench protection and refrigeration systems. For a twenty mile diameter machine there are many more instances where a distribution of controls becomes of considerable importance. A quick example is illuminating. The time of flight of a proton around the Doubler is on the order of twenty microseconds. Since the proton is essentially moving at the speed of light, that is also the time of propagation of an electronic signal around the accelerator. Now if someone at a central point wishes to request some piece of information, that implies some sort of round-trip time. (An outbound request for information and an inbound return of data.) This will be true unless a decision was made in advance that the information would always be wanted, so that arrangements were already made to have the data flowing inbound from its source without individual requests. If there is only one serial communication systems (as in the Doubler), it is only possible to make some 50,000 four-mile round-trip communications per second unless those communications are already very complex or, as noted, one-way communications have been previously established and the data is constantly flowing. If the dimensions of the accelerator are expanded by a factor of twenty, then the time of flight of protons around the accelerator is on the order of 400 microseconds and the number of electronic round trips has been correspondingly reduced. The number of requests for information on demand which can be accommodated has been reduced to about 2500 per second. That is a very low number and it absolutely implies that there will be a fair amount of distributed local control. Even intermediate control decisions are going to have to be made in a distributed fashion around the SSC ring.

The completion date for the SSC was given as possibly 1994. It is now about halfway between 1994 and 1974 when some of the early ideas were first considered for the control system for the superconducting Doubler. The differences in fundamental electronic technology which have developed between 1974 and 1984 are probably substantially greater than what one can expect to see in the next decade. Although the rate of development of electronic capabilities will probably continue to accelerate, many of the basic devices that one might need for a large, distributed control system ten years from now are available at present. Undoubtedly these devices will be subject to considerable improvement by 1994. But in 1974 they were not available. The first commercial microprocessors were announced in 1975. The planning documents from 1974 for the Doubler are quite interesting because those involved did not consider the distribution of computer control at all in the way it is now used. This is an example of a "fundamental change" that one does not necessarily expect to see every decade.

The Doubler Control System

As seen in Fig. 1, the Doubler has a basic centralized computer system which consists of two Digital Equipment Corporation VAX 11/780's and about twenty-one DEC PDP-11's which are networked together using DEC-PCL (parallel communications link) hardware. There are undoubtedly aspects of this particular choice of hardware which can be improved. Independent of the location of the SSC these improvements may be carried out at Fermilab during the course of proton-antiproton collider experimentation. Fourteen of the PDP-11/34's shown in Fig. 1 support consoles for operators. Each console is identical. The computers support the consoles on a one-on-one basis. This has proven to be a very useful and friendly sort of organization. It has been accepted as the type of organization that one would continue to support for the SSC. However, the use of something on the order of a DEC microVAX, or the equivalent from another vendor, but something which has considerably more computational power than a PDP-11/34 is probably indicated. The PDP-11/34's drive some hardware (not shown in Fig. 1, but indicated in Fig. 2) which permits serial communication to the actual console location, possibly over some distance.

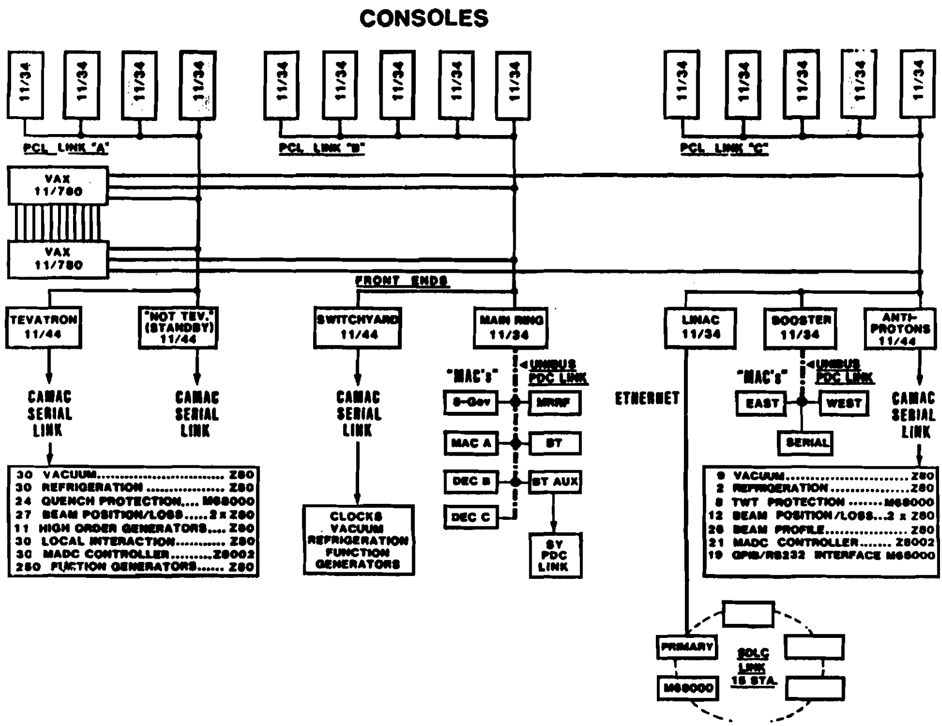

Fig. 1. The Fermilab accelerator controls system "ACNET."

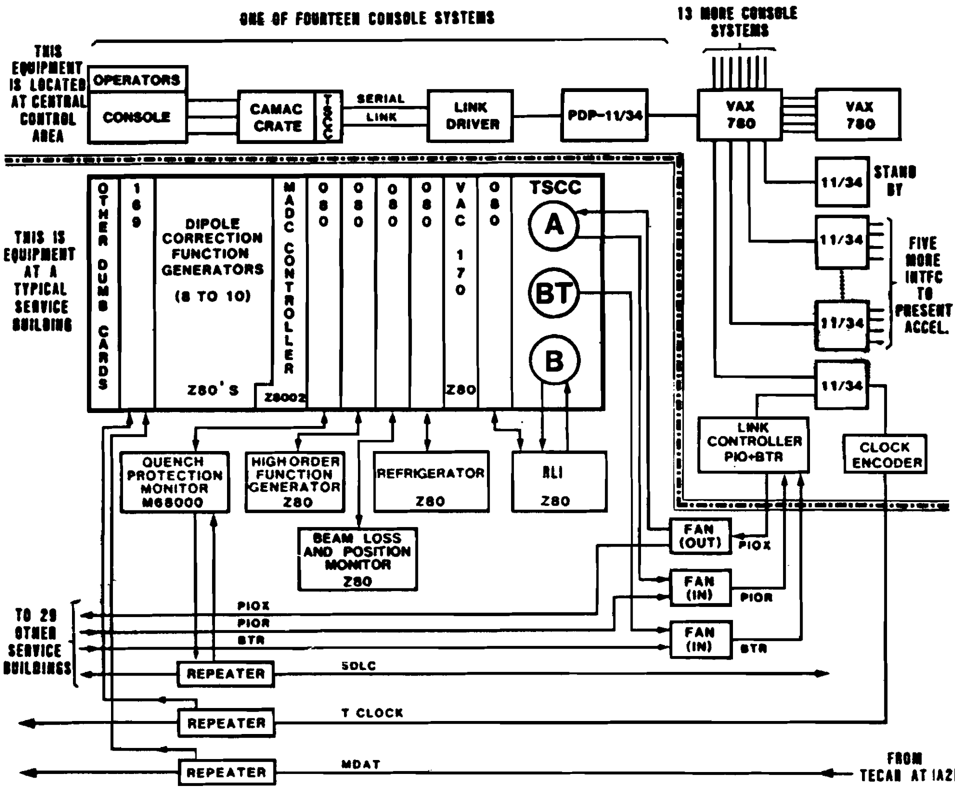

Fig. 2. The Fermilab Doubler controls systems architecture.

For the Doubler it has therefore already been necessary to address the question of console control of the accelerator at distances of several miles from the central control location. The present approach outlined above permits serial communications between the console equipment and its centrally located PDP-11 computer. In the future, this is almost certainly not a reasonable communications design when the distances are not two miles but rather twenty or forty miles. Rather, a greater utilization of long distance computer networking, with the console computers distributed to remote sites in addition to the console equipment, will almost surely be required. This will permit faster local access to data and also will greatly reduce the number of serial communication lines.

In many respects the Fermilab accelerator complex is similar to what is labeled the SSC injector; however, if an SSC is built from scratch in the four corners country of Colorado and New Mexico, the builders will not be faced with some of the historical imperatives that drove the controls group at Fermilab to do some very strange things when preparing to control the Doubler. Once some equipment is in place, it is seldom possible to replace it all at one time. There were some instances in the older systems of the conventional accelerator where very substantial changes were made in order to get a more unified control system going for the Doubler.

There are about 700 microcomputers distributed around the Doubler and Antiproton Source. They generally come in multiples of the service buildings around the rings. There are thirty vacuum controllers, thirty refrigeration systems (each one of which is driving about eleven closed loops), and there is a quench protection system which uses a somewhat more powerful microprocessor (the M68000) as compared tothe Z80 microprocessor used in the other systems. This quench protection system also has an independent link all around the Doubler ring because it is a system which is very dependent upon knowledge about conditions in immediate neighbor houses as well as upon some information from all the other houses. It is a rudimentary example of what might be called a local area network which comprises the entire Doubler ring so that any one quench protection system can query any other. There are beam position and loss monitors, and about 250 dipole correction function generators. Each one of the correction coil packages is driven by an independent function generator. This is an example of a system used both for the ease of human understanding of the correction functions and also installed for reliability considerations. It would have been possible when this system was being designed to put about eight of the correction functions in a package supported by a single 68000 microprocessors instead of using eight independent Z80 microprocessors; however, it is actually possible to run the Doubler with some of the correction coils missing. As a matter of fact, they are missing in one instance because they were not installed, a mistake made during assembly. In about four other places they are missing because the correction coil packages have proven to be one of the physically weak links in the actual superconducting accelerator and four have been destructively damaged. However, if one cannot make a correction at one place, it is possible to rework the the entire ring-wide correction function as distributed around the accelerator to take care of the problem. The extreme independence of this system makes it possible to either lose a single microcomputer (which seldom has happened) or to delete a physical correction element fairly directly. On the other hand, the loss of eight neighboring correction elements, due to the failure of a multiplexed system, would be more serious.

There are several examples of local area networks in the Fermilab accelerator complex. The quench protection system has already been mentioned. The system used in the re-instrumentation of the injection linac is another. In general, a local area network uses a straightforward technology like Ethernet, or a token passing Ethernet, or more generally, any number of computers requesting information from each other in an arbitrary fashion using any networking protocol which is independent of the details of the computers involved. The linac system used an SDLC protocol which is an IBM pre-Ethernet system characterized by a circular serial transmission of messages. A small segment of Ethernet is actually used to connect the linac system to the central host network.

The SSC Control System

As noted earlier, the physical size (rather than the complexity) of the SSC probably indicates that there will be considerably greater utilization of local area networks than was the case at Fermilab to date.

There are other features of an SSC characterized by only twelve major refrigeration centers and access points which are different when compared to the Doubler. At the Doubler, there are 24 service buildings servicing the six 60 degree arcs of the accelerator and six service buildings that are controlling the six straight sections for a total of about thirty service buildings around the four mile ring. At the moment, all of the major electronics is upstairs in the service buildings. In other words, the electronics is out of the tunnel where it is accessible. At the SSC there will be several kilometers between access points and therefore a considerable distance between an access point and an arbitrary controllable device. In such a configuration, the cabling cost quickly becomes rather expensive if all signals were to be brought back to the twelve access points. As a result, people have decided to study the question of actually distributing the electronics in alcoves in the tunnel of the SSC. This has several implications. The tunnel environment will probably be more hostile than the environment found in the Fermilab tunnels. The SSC tunnels are likely to be somewhat damper (since there will be no conventionally powered iron magnets to warm the air) and, of course, it will not be possible to get at the electronics to service it during acceleration or experimentation. Reliability, redundance, and backup will become important considerations. These are all related to the question of "What happens when something doesn't work right?" An important input for these considerations will be dependent upon guesses as to how the SSC will be operated. The operating scenarios will have to come from the accelerator system designers, the accelerator theorists, and the experimental physicists. One scenario might be that a proton-proton or a proton-antiproton fill would be done once per day and that experimenters would be reasonably happy if they got twenty hours of colliding interactions and then had from two to four hours of access to the tunnel for various purposes. At the Doubler, access to the tunnel is not a fast thing to arrange. One basic problem is the exposed electrical buswork in the tunnel. At present, it is a fairly lengthy procedure simply to get the breakers undone in order to allow general access into the Doubler tunnel. Another problem is the oxygen deficiency hazard. This sort of detail must be considered if the scenario includes a plan to service electronics inside the tunnel on a daily basis.

To summarize, there probably will be electronics in the tunnel of the SSC, there probably will be an emphasis on local area networks, and there probably will be an increasing emphasis on reliability. Note that low voltage digital and analog controls are by no means the whole story when considering accelerator reliability. What fails in general is not microcomputers and integrated circuit chips. One recent microcomputer "failure" at the Doubler occurred when the service building roof sprang a leak and a rainstorm drenched the microcomputer. This is an example of the "real" problems a person has to include in one's thoughts.

The SSC accelerator itself will not work if a single major bending dipole is "missing" (out of the circuit). A relatively small loss of bending angle at the one point where a dipole fails will result in substantially less than one centimeter of orbit displacement at the point of failure. However, a "bump" is thereby put into the orbit that results in a ten centimeter displacement at some other point around the accelerator. That would be well outside the aperture of the proposed magnet system. This means that there are certain situations where it is not possible to protect against catastrophic failures.

During the early stages of the SSC preliminary design, some people thought that they might be able to build a truly "passive" quench protection system into the magnets themselves. The SSC reference design does not make that assumption. Magnet builders also feel that it is probably not possible to do so. This means that a very "active" electronic quench protection system for magnet protection will continue to be needed. This was, one recalls, one of the important arguments for the distributed processing utilized in the present Doubler system. At the SSC, the same emphasis only becomes greater. There simply is not time to collect all of the necessary information for evaluating a magnet's superconducting status, and then deliver it to, and process it at, a central location. Typically, a decision as to whether a magnet has gone "normal" has to be made in one or two 60Hz line cycles. At that point, something must be done or the physical integrity of the magnet is at risk. The quench protection system at the SSC will face timing constraints identical to those at the Doubler.

This means that it will be necessary to have some relatively high powered processing capabilities in the distributed locations at the SSC. Some of the SSC reference designers have argued that it may be possible to combine many of the functions, which are separated into individual microcomputers at the Doubler, into larger microVAX or equivalent systems. This is certainly possible but whether it is completely desirable requires some substantial program or failure analysis to evaluate. The SSC reference design, for example, suggests that the same rather high powered miniVAX that would do quench protection of a section of arc would also act as a communications node at a major access point in the network. This could place conflicting demands for processor cycles in a fashion that would be irritating to one or the other of the processes, say quench protection or communication.

One of the features that is very nice for Doubler accelerator operators is the ability to connect a "trackball" (similar to a computer mouse but a little different) directly into the system in real time. This permits one to adjust the numerical value of a variable, have the altered value be sent to the hardware, and then have a new reading of the variable be returned from the hardware at 15Hz. This is fast enough so that from the human (physiological) point of view, there appears to be a real time physical connection. Earlier in our discussion it was noted that the total number of round-trip communications between a remote location at the SSC and some central facility will be reduced to only 2500 or so per second. If, in addition, there were to be too many "layers" of local networks and interfacing computers between the local area networks, there could begin to be some difficulties in trying to pass information along in real time in order to provide some type of response for humans approximating 15Hz.

For the Doubler, the data base for the entire system is centralized in the "Operational VAX." There were some strong arguments in favor of this approach from the people that proposed this, advocated it, and implemented it. With a widely distributed system as at the SSC, it may be desirable to turn on the refrigeration systems, for example, as arcs of the SSC are completed. If the complete central control system is not done at that time, a greater distribution of the data base may be desirable compared to what has been done for the Doubler. However, there are certain problems that arise which are among the reasons that one did not choose to distribute the data base in the Doubler system. The primary information in the data base at the SSC would still be addressing information. In other words, the data base contains the necessary information so that if somebody at any console wants to get at a particular piece of information, one picks up a road map and this map tells one how to get to the information. The road map is handed off further down the line to all other computers involved until one reaches the computer with direct access to the information requested.

Now, one way to avoid the necessity for some of the complexity of the road mapping scheme is to fold into the device-naming-architecture a great deal more directive information than is implicit in a device name at present in the Doubler. There are many people who advocate the practice of a "meaningful mapping" of all device names, and there are some very strong reasons for wanting to do it, avoiding some of the reliance on a centralized mapping data base.

The next three figures illustrate the system proposed in the reference design for the SSC. It shows a system of super-mini's (for example VAX 11/780's or 785's), large mass storage, and twelve operator consoles. The super-mini's do not physically drive the consoles; the consoles are perhaps driven one on one by computers each equivalent to a microVAX. The super-mini's are networking systems and network switch controllers, with perhaps the additional job of providing a redundant system for verification. The SSC in the reference design uses a standard long haul network to distribute local control around the 100 kilometer ring. Figure 3 illustrates a node on a local area network for one of the twelve sectors of the ring, or perhaps for the injector complex. Figure 4 illustrates the proposed architecture. The long haul network is shown, as well as an array of mini computers and some subsidiary local area networks. Microprocessors are located inside the tunnel. Each microprocessor is illustrated managing up to ten "half cells" of magnets. There are switches to permit some redundant paths for communication. Figure 5 shows the architecture at the cell level. This is shown including a 16 bit microprocessor that is managing a number of subsidiary modules in an interface crate. This is the unit that would be repeated most frequently in the tunnel.

The reference design proposal is certainly an example of a system that could be built today. From the technology point of view, there is no overwhelming difficulty with the design. It would probably produce most of the features that have been found necessary in the operation of the Doubler to date, with the possible exception that there might be a little difficulty in making a real time connection between an operator's control device and a piece of accelerator equipment with a 15Hz response.

At the SPS at CERN they do not have a 15Hz operator-to-device connection capability. They do not try to give operators a real time feel of control over particular devices. People who have played with both systems, however, feel that the Doubler system has something to be said for it.

Notice that most of the subsystems mentioned as part of the Doubler are shown in the SSC reference design figures. These include beam position monitors, beam loss monitors, wave form generators, and voltage monitors. The design report does not discuss whether these should involve subsidiary microcomputers or

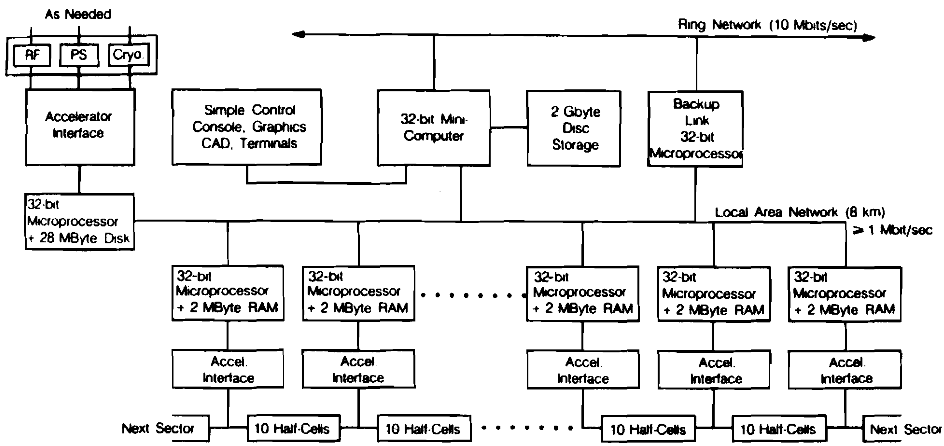

Fig. 3. SSC sector control system block diagram.

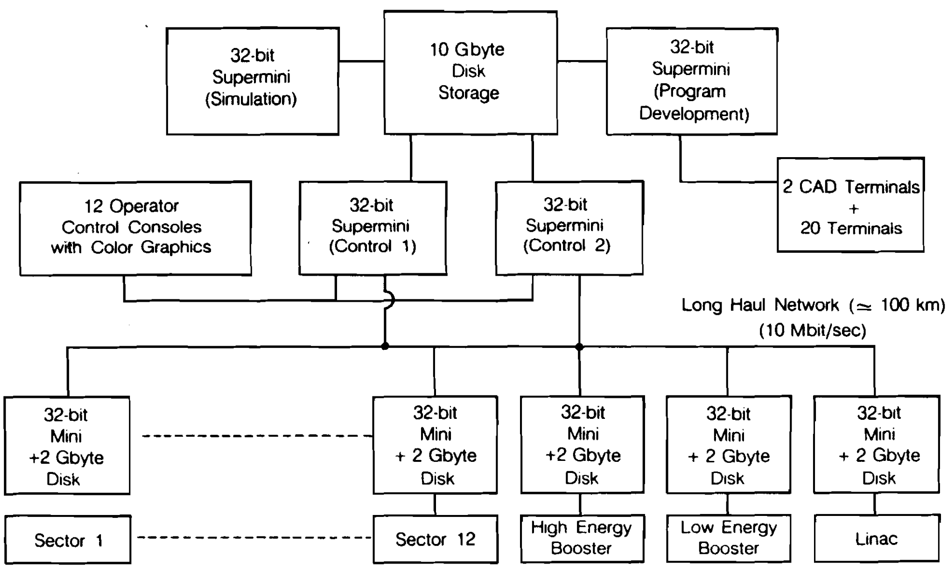

Fig. 4. SSC central control system (architecture overview).

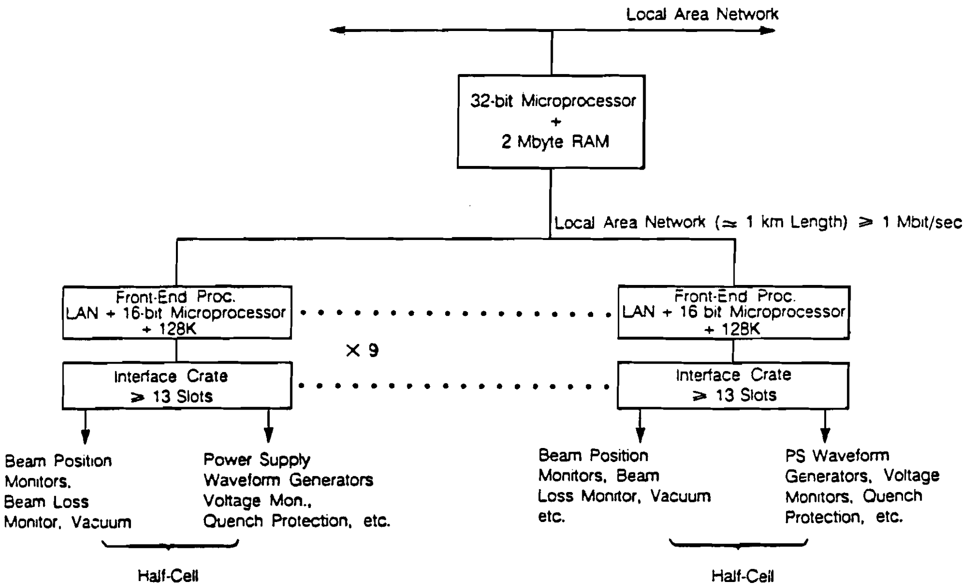

Fig. 5. SSC distributed microprocessors.

be multiplexed out of a single high powered 16 bit microcomputer. That is an example of something the final details of the design would have to address. The similarities to the Doubler system are obviously very large.

Summary

There are a number of challenges with respect to subjects such as the requirements for local area networks, the distribution of the data base, the question of the local distribution of electronic equipment in the SSC tunnel, reliability, and redundancy. Fundamentally, the control system is not a system that will have a significant impact on the basic questions concerning the possibility of constructing the SSC. It is certainly correct to place the major part of the research effort into developing the magnets as well as defining the actual physics goals of the experimentation. The control system can undoubtedly be matched to the requirements so defined with equipment available today, and certainly with equipment to be developed over the next three or four years during the R&D phase.