Controls for the CERN Large Hadron Collider (LHC)

K.H. Kissler, F. Perriollat, M. Rabany and G. Shering
CERN, CH-1211 Geneva 23

Abstract

CERN's planned large superconducting collider project presents several new challenges to the Control System. These are discussed along with current thinking as to how they can be met. The high field superconducting magnets are subject to "persistent currents" which will require real time measurements and control using a mathematical model on a 2-10 second time interval. This may be realised using direct links, multiplexed using TDM, between the field equipment and central servers. Quench control and avoidance will make new demands on speed of response, reliability and surveillance. The integration of large quantities of industrially controlled equipment will be important. Much of the controls will be in common with LEP so a seamless integration of LHC and LEP controls will be sought. A very large amount of new high-tech equipment will have to be tested, assembled and installed in the LEP tunnel in a short time. The manpower and cost constraints will be much tighter than previously. New approaches will have to be found to solve many of these problems, with the additional constraint of integrating them into an existing framework.

I. LHC REQUIREMENTS

The Large Hadron Collider (LHC) is the major project planned by CERN[1], and will be its largest and most expensive ever. It will present control problems much greater than those experienced in earlier accelerators.

LHC is a superconducting twin beam hadron (proton initially) collider providing 7.7 TeV per beam at 10 Tesla bending field. The novel twin bore magnets in their cryostats will be installed in the same 27 kilometre tunnel as the LEP machine. The **scale** of the control problem can be gauged in part from the 1792 dipole and 392 quadrupole cryostats filling most of the circumference, in part from the number, about 2000, of insertion and corrector magnets and appropriate beam instrumentation. The **difficulty** of the control problem will come from the sensitivity of the superconducting magnets to quenching under beam loss from the 4725 bunches of 10^{11} protons at 400.8MHz, making 851mA. This problem is exacerbated by the time varying persistent currents and the need for strong collision insertions to achieve the targeted luminosity of over 10^{34}. These requirements will strain dynamic aperture and magnetic field control to the very limit.

As LHC will be built in the same tunnel as LEP, a lot of equipment and controls will be common to the two machines. A major objective will therefore be a seamless integration of LHC and LEP controls. This will not be easy, in part due to the much more difficult control problems of LHC, in part due to the wide separation in time between the construction of the two systems compared to the speed of evolution of controls technology. A challenge for LHC will therefore be to permit the use of the latest and cheapest controls technology in such a way that it integrates with existing technology, allows experience and algorithms to be maintained, and does not demand a difficult and costly upgrade of existing systems. Another objective to be borne in mind is the aim of having a single control centre for the whole of CERN on the LHC time scale.

II. MAGNETIC FIELD CONTROL

This will be the most difficult control problem for LHC. The magnetic field is determined not only by the voltages and currents from the 1400 power supplies, but also by "persistent currents" in the superconductor which vary with time depending on the history of the magnetic cycle. Fortunately HERA experience has shown these effects to be reproducible, hence eventually calculable and correctable. Corrections will be derived from magnet measurements during the construction and from on-line measurements from reference magnets. The final trimming will have to be done using single pilot bunches of first 10^9 then 10^{11} protons. After full beam injection, continuous feedback control will be required, especially immediately after injection and during beam squeezing.

A. Modelling Server

The solution envisaged is to use a modelling server which will re-calculate the power supply settings in real-time, using measurements from the reference magnets, beam measurements, past history of the magnetic cycle, and the magnet characteristic data-base. The update time will be between 2 and 10 seconds. Tests on an Apollo DN10'000 in the LEP control system indicate that the computational load will not be beyond the sort of on-line computer we can expect in the LHC time scale. Each power supply will have a microprocessor capable of interpolating the required voltages and currents between modelling server updates.

B. Fast Communications

A new communications system is being studied for LHC, in conjunction with SSC, in order to acquire the beam data and set the power supplies at the required rate[2]. This will use TDM communications technology and reflective memory computer cards to provide parallel transmission of data with no software protocol overhead. The power converters and beam monitors will use the latest digital signal processors and microprocessors, so no problem is envisaged at this level provided they can all operate in parallel.

C. Pilot Procedures.

It will be necessary to automate the procedures of machine preparation, pilot injection tests, and full beam injection and acceleration, in part to keep the persistent current effects under control, in part to avoid human error. Much work of this kind was done for the Antiproton Accumulator after initial commissioning, but for LHC this will have to be done beforehand, so placing an early load on the applications programming effort. Of course the modelling programs in the servers will also have to be ready and tested before injection tests can start. High quality work in these areas will greatly reduce the time needed to get the machine into proper working order.

III. QUENCH PROTECTION

The first line of quench protection will be assured by 24 quench protection stations distributed around the ring. They will constantly analyse the voltage transducers connected at each coil access point in order to determine if a quench has taken place. The quench station will then take a series of time critical actions including firing the beam dump and firing the magnet heaters to spread the quench evenly over the magnet. This forces the current to flow through the cold diode protection shunts to avoid local damage to the coil at the quench position. Helium pressure relief valves are then opened to avoid reaching emergency over pressure. A 5 second pre-quench record will be held for "post-mortem" analysis.

The magnet control computer will then take over, switching in series resistors in appropriate places to run the current down as quickly and safely as possible. Constant surveillance of the quench protection stations will of course be necessary as no malfunction can be permitted. Also the state of each magnet will have to be monitored for abnormalities in temperatures and voltage drops to detect deviations which might indicate variations in performance from the model, or incipient trouble. Before each injection extensive automatic checks will have to be done on all magnets, cryogenics and quench protection stations to ensure a fully normal situation.

All of this software will have to be ready and well tested before any significant number of magnets can be put into operation. For this reason the controls effort will liase closely with the magnet test strings and test stands right from the beginning.

IV. BEAM MONITORING

The two main subsystems are the orbit measurement and the beam loss monitors. The monitors will be wired to 24 concentrators round the ring for local processing and connection to the network. Each local unit will have a direct connection to the beam dump which will be activated if conditions which will lead to a quench are detected. In the case of some critical beam loss monitors, like those near the collimators, the dump will have to be activated in 1 or 2 revolutions, ie. 90 to 180 microseconds.

The closed orbit measurement will be systematically acquired by the modelling server to aid in its 2-10 second update of the power supply currents. In addition to helping compensate for the persistent currents, this will provide a continuous on-line correction of the orbit.

The orbit and beam loss monitors will also be acquired on a regular basis by the central alarm server to help in the detection and avoidance of quench provoking situations.

Other instrumentation is foreseen to measure tune, chromaticity, profile, and dynamic aperture and will also be controlled by local sub-systems. These measurements will be used along with the orbit and loss monitors by the programs which set up the machine for injection, acceleration, and collision. A particularly important role for all the instrumentation will be the pilot injection tests first with a bunch of 10^9 protons, then with a 10^{11} bunch to prepare for the full batch.

The beam monitoring stations will be connected by the fast acquisition system to a beam data server and hence to the modelling server. A first trial of the fast acquisition is planned in the near future using some of the LEP beam instrumentation as a prototype.

V. CRYOGENICS

The number of pieces of cryogenic equipment to be controlled is very large. However, the equipment involved is well known to industry, and an industrial control system will be purchased for this purpose. This must communicate effectively with the main control system, as any magnet whose cryogenic state or pre-history is inadequate may be prone to a quench before full field is reached. Thus the state of each magnet must be checked by the control system before each cycle and continuously monitored thereafter. Also access to the cryogenics systems should be the same for normal operator use as the access to any other system.

VI. OTHER SYSTEMS

At present it is assumed that most other LHC subsystems will be controlled by extensions of the LEP control system. If after further study it is found that some other system, for example the collimators, needs special fast control, they can be dealt with using the fast techniques used for the beam monitors and power supplies.

At first LEP and LHC will not function at the same time and therefore much of the standard LEP controls may be re-used for LHC. In any upgrade of LEP this will be borne in mind and spare capacity will be provided where appropriate.

VII. LHC DATABASE

A database (ORACLE) will be used to store the characteristics of the LHC machine and to help with the planning and installation. Such a facility proved invaluable for LEP. For LHC there will be even more high technology equipment to be installed in the tunnel in a shorter time and with limited staff. A maximum amount of easy to use informatics must be provided to ease and control this installation.

All controls information will also be stored in a database. The use of an on-line database was pioneered at the PS on the original IBM 1800, followed by the data-module concept at the SPS. The usefulness of having cabling and installation data on-line was illustrated in the '70's with the SPS experimental area system[3]. For LEP it was necessary to copy the data to local files for use in the control system. We hope to avoid this in LHC and use direct on-line data-base access. Not only will this save work and additional software, but it will reduce the possibilities of errors and inconsistencies. Two approaches are envisaged. For interactive use the database can be interrogated as required by SQL commands during the course of the applications programs. Tests at the PS have shown that response times consistent with human operator interaction can now be achieved. Alternatively, applications programs which must run fast or which need a lot of data, can interrogate the database at start up time and store the data in RAM structures for immediate and fast use, as is done at Isolde[4]. This also protects the programs from inconsistent updates. For any updates to become effective, the program has to be stopped and restarted.

VIII. ARCHITECTURE

Controls technology sees major revolutions on something like a five year time cycle. As LHC completion is more than 5 years away it is premature to make final decisions on the architecture, even more premature for the detailed components to be used. Nevertheless, since this is a controls conference, an architecture is presented for discussion which illustrates much of the current thinking[7]. This is shown in figure 1.

A. Equipment Connection Policy

A major feature is the equipment connection. An important change in equipment connection policy and technology took place in LEP. The responsibility of the equipment layer of the controls was transferred to the groups responsible for the equipment, and a "thin" connection made to the control system. This policy will be continued and reinforced for LHC . The controls group will provide the hardware connection for data and timing, and together with the equipment group it will define a control protocol which will define the characteristics of the equipment to the control system. This policy has several impor-

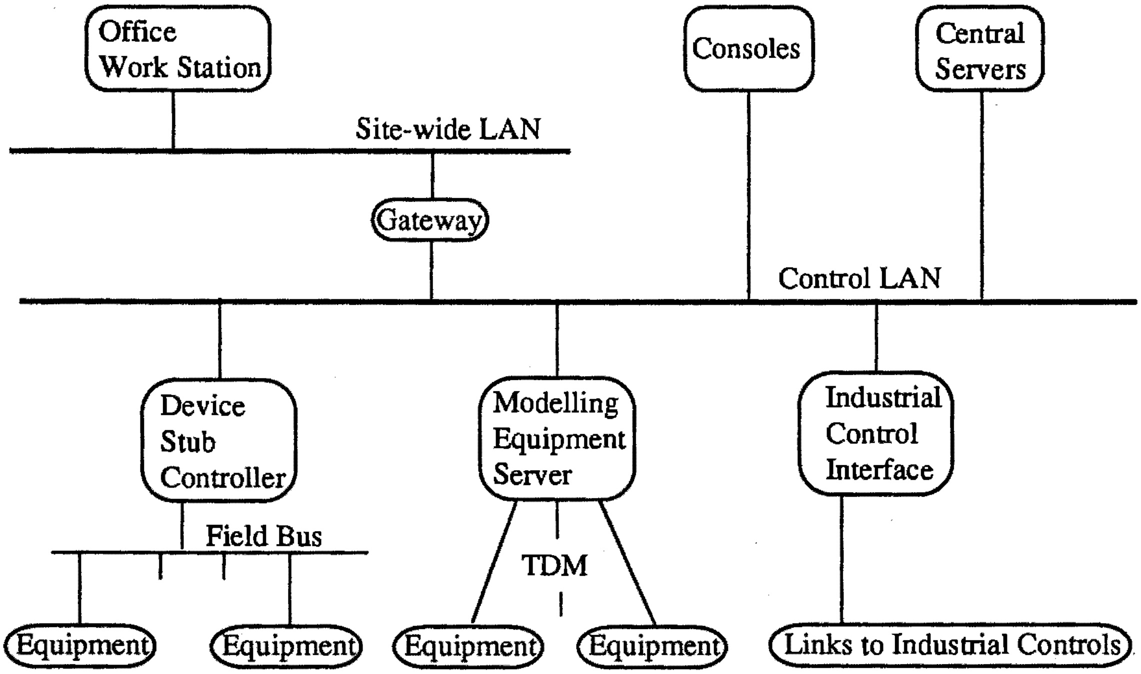

Figure 1. Schematic diagram of architecture

ISBN: 978-3-95450-254-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS1991-S03SRD01

tant advantages. Firstly it clearly defines the budget separation between the control and equipment groups. Secondly it clearly delineates the responsibility for design, performance and maintenance of the control equipment. Thirdly it decouples the evolution of the control system from the evolution of the specific equipment controls. This is becoming an ever more important consideration as the park of control equipment increases and we have not the resources to change everything at the same time.

B. LEP Equipment Connection

The bottom left of Figure 1 shows the LEP equipment connection[5]. The specific equipment is connected to a field bus, MIL 1553, which is connected to a "Device Stub Controller" consisting of an Olivetti PC running the SCO UNIX operating system (replacement by LynxOS, a real time POSIX compliant operating system is planned in the near future. This operating system can also run in a VME crate). These PCs are located in the field, and are linked to the rest of the control system by the control LAN. Much of the LHC equipment will be in common with that of LEP. This includes, of course, all the tunnel management services. Some of the new equipment of LHC may also be connected to the MIL 1553 to save money and development cost. As LEP and LHC will not operate at the same time, economies can be made in sharing equipment and resources.

C. Fast Equipment Connection

The equipment connection shown in the centre bottom of figure 1 is a new development for LHC, being done in conjunction with the SSC[2]. The driving force for this is a fast connection to allow the rapid measurements and power converter updates required for control of the magnetic field. Other advantages of low cost and simplicity are also claimed for this new method, however, which may make it a contender for applications which do not need its speed. The basis of this method is to use Time Division Multiplexing (TDM) techniques to provide an individual channel from each equipment to a central server, all connections working in parallel with no software protocol overhead. This TDM technology is in commercial use for telephone and higher speed data-link requirements by the communications industry, mostly running on optical fibre. It will also be used in LHC for dedicated connections, telephone and video links, interlocks, and many other applications which are traditionally implemented using hard wired copper. Between the equipment and server, the TDM channel will link "reflective memory" cards in which part of the equipment memory containing the control protocol is reflected into the central equipment modelling server. Thus the modelling server can read its input, do its calculations, and write its output with no degradation in performance due to software protocol overheads and long transmission delays. Several servers may be used, eg separate ones for beam monitors, power converters and magnet monitoring. The TDM may be used to link these servers among themselves using reflective memory, or to link the equipment to several servers.

D. Industrial Equipment Connection

The third mode of equipment connection, shown diagrammatically on the bottom right of figure 1, is connection through industrial control equipment[6]. For the cryogenics equipment, and perhaps others, it makes sense to buy the controls from the equipment manufacturer or an industrial process control manufacturer as he has solved the problems, has all the components in house, and can take charge of the installation and maintenance. The controls thus purchased may stop at the level of the Programmable Logic Controller (PLC) for the equipment concerned, or even a complete system may be bought with the manufacturer's network system and consoles so that he can take full responsibility. In both cases an interface to this manufacturer's equipment has to be provided to achieve two way communication and control, and to integrate the industrial controls into the protocols and methods of the rest of the controls.

E. The Central Servers

In addition to the modelling server connected to the power supplies, there will be other central equipment servers connected by TDM to subsystems requiring fast response such as beam monitors, quench protection stations, collimator subsystem controls, dump, etc. Other central utility servers, as shown on the top right of figure 1 will be required for alarms, a file server, an on-line database engine, etc.

F. The Consoles

These will be standard large screen workstations. The PS and SPS console facilities are being upgraded to match LEP[7]. A considerable effort is being put into this and we can assume that the resulting facilities will be entirely adequate for LHC. There is a strong desire to have a single control centre with standardized facilities for the whole of CERN on the LHC time scale.

G. The Office Workstations

LHC will involve a large fraction of CERN's accelerator community over the construction and commissioning period. These people will be spread over a large geographical area. To permit them to work without too much travel to the central control room, access will be provided from the office workstations connected to the site wide LAN. A gateway will be used to authenticate access permissions, and to cut off access at critical times.

H. The Control LAN

A single LAN is shown diagrammatically in figure 1. The present system uses both Ethernet and Token Ring in a bridged configuration. It is reasonable to expect the use of FDDI in the LHC time scale.

IX. PROTECTION

LHC will have severe requirements concerning the authorization of actions through the computer system. In situations where a quench is possible, equipment changes must be restricted not only to authorized persons but also to authorized programs which have been vetted to perform adequate checks for quench avoidance.

There will be a large number of access points to the system for testing, maintenance and commissioning including, as mentioned above, the office workstations. These numbers will aggravate the protection problem. Network management will have to provide access permission or denial facilities, as now being added to LEP, and the alarm and surveillance system will have to detect and locate abnormal access attempts. The protection will have to be closely linked to the machine operational state. Widespread access will be needed to speed installation, testing and repairs, moving swiftly to a closely controlled situation when quenching becomes possible.

X. TIMING

In addition to data transfer, many types of equipment need special timing signals and events. General purpose timing systems have been developed for PS and SPS from different historical backgrounds. Currently work is progressing on new designs which will cover both areas and which will take into account LHC requirements. Thus the timing distribution for LHC should be part of a new overall system for all CERN accelerators. There is a hope that the timing information may even be combined with the data stream so reducing the number of cables, connectors, and chips, so reducing cost and increasing reliability.

XI. APPLICATIONS PROGRAMS

The maxim "a control system is only as good as its application software" is as true now as ever it was. In the '60s the term "software barrier" was coined to express the difficulties in achieving good applications. In the '70s a determined effort using NODAL in the SPS, PS, PETRA and TRISTAN achieved a good measure of success. In the '80s the problem re-appeared with more complex systems and higher demands. For LHC a serious effort will be made to put together a team of machine physicists and programmers who will ensure that the applications programs will be available, not only to run the machine but also to make a positive contribution to a fast, smooth and safe commissioning.

For LHC, applications software will be required on two levels, at the server level, and at the console level. In the servers sophisticated software will be required for the modelling and real time closed loop control of the high intensity beam in the superconducting magnets. High reliability will be required in this software, and in the quench protection and magnet surveillance programs. All changes and machine settings will have to be verified through a model before being applied so as to avoid quenches.

In the consoles a wide range of software will be required. Automated injection procedures will be needed for speed and reliability. These programs and the server software will have to be ready before commissioning, unlike the cut and try approach which could be used with earlier less sensitive accelerators. Extensive software for beam measurements and diagnostics will be required.

Surveillance and alarm programs will be particularly important to avoid quenches and warn of quench provoking situations. A variant of this software will be required to check out the machine before injection, or before any action which might result in a quench.

Some of this software will have to be professionally written, installed and tested, especially in the servers. Other parts will be better written by the machine physicists and hardware specialists themselves. All will have to be carefully specified beforehand. Tools will include professional development facilities for compiled programs, the old tried and tested NODAL[8], mass market packages such as spreadsheets linked to the machine variables as used at Berkely[9] and in Isolde[4], and other commercial synoptic packages for control as used in LEP.

XII. REFERENCES

[1] Design Study of the Large Hadron Collider (LHC), CERN 91-03.

[2] C.R.C.B. Parker, M.E. Allen, H.C. Lue and C.G. Saltmarsh, A New Data Communications Strategy for Accelerator Control Systems, Nuclear Instruments and Methods in Physics Research, 1990 (Proceedings of Vancouver Controls Conference, 1989).

[3] H.W. Atherton, M. Rabany and R. Saban, A database as a bridge between hardware and software, Proceedings of the 5th International IASTED Symposium, Tunis 1982

[4] R. Billinge, I. Deloose, A. Pace and G. Shering, A PC Based Control system for the CERN ISOLDE Separators, Proceedings of this conference.

[5] P.G. Innocenti, The LEP Control System, Nuclear Instruments and Methods in Physics Research, 1990 (Proceedings ofVancouver Controls Conference, 1989).

[6] M. Rabany, Interfacing Industrial Process Control Systems to LEP/LHC, Proceedings of this conference

[7] R. Rausch and C. Serre, Common Control System for the CERN PS and SPS accelerators, Proceedings of this conference.

[8] M.C. Crowley-Milling and G. Shering, The NODAL system for the SPS, CERN 78-07.

[9] S. Magyary, M. Chin, C. Cork, M. Fahmie, H. Lancaster, P. Molinari, A. Ritchie, A. Robb, C. Timossi and J. Young, The Advanced Light Source, Nuclear Instruments and Methods in Physics Research, A293, 1990.