14

Control System

Contents

14.1 Introduction

Control system implementation is notoriously dependent on available technology. "Notoriously dependent," because its technology base—computers, communications, and electronics—is developing at the most rapid rate of all components used in accelerators. Thus design proposals for systems to be built in five to ten years from the date of a proposal are often little more than examples of how dated ideas can become in that period of time. In spite of that fact, some rough cut on control system design must be made now. The control system will interact with most other components of the accelerator. It is in everyone's interest to define the most important features—in terms of functionality, not implementation—of hardware and software interfaces early. This will promote uniformity where appropriate and will give other systems an opportunity for review.

In this chapter, we will acknowledge that proposals for a control system design are transient and focus on two points:

1. Although the implementation design of a control system may be transient, the functional requirements that such a system must meet are much less so. Operations at the SLC and the Final Focus Test Beam have given us valuable lessons for the control system needs of an NLC [Humphrey 1992]. Thus, the first section of this chapter will describe some of the key functional requirements for NLC operations that the control system must supply and the second section will discuss some of the implications of these requirements for the architecture.
2. A serious proposal to build NLC must include a cost estimate for all significant components. The cost estimate for the control system can only be based on a specific implementation; therefore, an implementation model, based on presently available technology, is proposed in the third section of this chapter.

14.2 NLC Requirements

A control system should be tailored to its users and the challenge they face. Much of its bulk is devoted to device controller interfaces, networks, consoles, etc.; those parts will not be discussed at length in this chapter, though some rough idea of the numbers of such devices will be needed for the cost estimate and for discussions of reliability (Chapter 17). But, there are many aspects of a linear collider that impose requirements on the control system that are distinct from normal accelerator control systems. Most of our understanding of these requirements comes from experience with the SLC and its control system, which has evolved considerably from its initial implementation. They include:

- Beam-based feedback
- Measurement synchronization with beam pulses and high-bandwidth data acquisition
- Special handling of damping ring applications
- Automated procedures for tuning
- Machine modeling and simulation
- Accelerator/detector coupling.

14.2.1 Feedback

Beam-based feedback, of the type described in Appendix D. should be considered a possible engineering solution for all tight tolerances which are static enough to be within the correction bandwidth of the feedback system. Furthermore, beam-based feedback should be considered for all systems whose initial fabrication errors place them outside of predicted tolerances. However, for systems with disturbances outside of the feedback bandwidth, or "within the bunch train," other solutions such as feedforward or dedicated, very high-bandwidth feedback systems should be considered.

There are several aspects to this overall approach:

- Use of the control system, where effective, to offset the cost of meeting tight tolerances by modifying or rebuilding the system affected.
- Development of feedback to control instabilities of all time scales, especially mid to long term. In the SLC "mid term" is the highest frequency presently addressed by the fast feedback system—120 Hz.
- Strong effort to counter fast instabilities with high-powered tools and feedforward.
- Consideration of high-bandwidth feedback for systems with very tight tolerances.
- Establishment of *treaty regions* with full separation of phase space and centroid stability.

It is critical to identify which technical issues can be addressed by these means and which cannot. Mechanical, thermal, slow magnetic fields, and low-bandwidth microphonic disturbances have time scales which can be addressed by such feedback systems. Pulsed devices such as kickers and modulators must be stabilized using other feedback techniques (not beam-based).

The present system of feedback at SLC was built after much of the lower-level design was mature. It evolved through several generations of development, each incorporating the lessons learned. The pervasiveness of the use of feedback throughout the accelerator controls was not fully anticipated and therefore the system is not optimally integrated. This requires attention in the detailed design of an NLC control system.

Fast Feedback Rate Considerations

The SLC feedback system is designed in a generalized, database-driven fashion, which contributes greatly to its flexibility. The system uses standard control-system hardware. As a result of this design, unplanned control loops can often be added with only database work, without requiring hardware or software changes. In addition to the beam position monitors and correctors used for steering control, the feedback system is capable of measuring and controlling a wide variety of devices. For example, it is as easy to add a steering loop in the linac as it is to provide control of laser gun timing in the injector. Special-purpose extensions to the linear feedback system have been added to accommodate non-linear cases, such as optimization feedbacks in which the measurement responds parabolically to actuator movement. The system also provides built-in diagnostic and analysis capabilities, and the many sample-only monitoring loops provide a wealth of diagnostic information. These design features have been key to the success of the SLC system, and the NLC design should be equally flexible and extensible in order to support unplanned controls needs.

It must be noted that for this generalized feedback system to perform properly sufficient resources must be available. The specific system used in the SLC employs dedicated point-to-point links in order to minimize communication overhead. A typical launch feedback loop at the end of the SLC Linac stabilizing angle and position in two planes for both electrons and positrons absorbs approximately 70% of the available bandwidth of an INTEL 486 single board computer running at a 66-MHz clock rate.

The feedbacks for the NLC are planned to run at the full beam rate of 120 or 180 Hz. The steering loops control the average trajectory of the bunch train rather than individual bunches, so the Q BPMs are used, which measure the average of the train. Dipole corrector magnets are used, for which the feedback controls the magnetic field by setting a DAC (digital to analog converter) which alters the current from a power supply. Correctors and other actuators need to respond (to make 90% of a requested change) in a single 180-Hz period.

For the NLC, linac feedbacks will use typically use 10 BPMs and four correctors (two horizontal and two vertical). Communications and CPU must support processing of this local data at the 180-Hz rate. In addition, a cascade correction system for the NLC will require 180-Hz communications over long distances between linac loops (in diagnostic sections).

14.2.2 Synchronized measurement

The collider will be a pulsed machine with a relatively low repetition rate. A linear collider does not have the tendency to stabilize inherent in a storage ring machine: every pulse is a new beam. Therefore it is crucial that the facility exist to take measurements synchronized to particular pulses throughout the accelerator. These will include beam measurments, hardware diagnostics, and state information from the detector. This will allow pulse-to-pulse correlations of all different aspects of the collider and is the only way to efficiently trace sources of beam jitter and subtle hardware failures. A full sample – data for every pulse – over an extended time may be needed to see both high and low frequency instabilities.

Pulse-oriented sampling is also needed to provide an estimate of electromagnetic fields and other analog monitors at the beam pulse frequency. This means that the acquisition of analog signal data should be synchronized with the beam or line frequency.

14.2.3 Data acquisition and processing

High data-acquisition bandwidth is required to characterize pulses for optimization or feedback. Ideally, the BPM system should produce phase space centroid information of each bunch in the train at the full 180-Hz rate. This will require a control system data acquisition that has a throughput equivalent to the maximum beam pulse rate. The system should be able to acquire and process data more or less indefinitely. With a large number of samples acquired at the pulse rate, fine details of machine performance can be examined such as, for example, the frequency structure of narrow excitation lines. (Such lines have been observed in the SLC beam motion frequency spectrum, and are sometimes caused by the linac water-cooling pumps).

A more serious challenge is the measurement of phase space volume or emittance. In order to measure the bunch train's transverse and longitudinal volume with the sampling mode currently in use at SLC, a minimum of several hundred pulses is required. At low repetition rate (as when trying to diagnose and correct a problem) this is too slow to allow for effective tune up. A profile monitor system has been proposed that can be used to “fast” scan the entire train in a single pulse.

14.2.4 Special architecture for damping ring applications

At SLC, most of the data acquisition system and scheduling system is optimized for a pulsed linac machine. A special architecture is required that can be used to track the performance of the damping ring hardware throughout a single storage cycle and can track the behavior of the beam at the same time. The B-factory control system, for example, will have to have an integrated timing system that handles both the injected and stored beam. NLC damping rings will need controls that address the problems of injection, extraction and damping.

14.2.5 Tuning

Complex tuning will be required in the NLC. This is true in many places but is most prevalent in the final focus. To limit the luminosity loss due to the tuning procedures, the tuning scans and measurements must be highly automated and fast. Thus, the entire scan, which usually involves changing the strength of many magnets and then measuring the IP spot sizes with a beam-beam deflection scan, would probably be performed locally before the data can be shipped back and processed.

14.2.6 Modeling and simulation

An early decision in the design of the SLC control system was to base many aspects of machine setup, operation, and diagnostics on online accelerator models rather than a look and adjust method of interaction. The adoption of this approach has resulted in the development of a rich suite of applications that forms the foundation for near-automated operation of SLC. This modeling framework has been effectively used throughout the machine life cycle including commissioning, routine operation, diagnostics, and performance upgrade and optimization phases.

Given the complexity, strict tolerances, and the expansiveness of NLC, it is vitally important to have an online modeling environment with appropriate degrees of sophistication to facilitate machine commissioning and operation. At SLC offline simulation has been used to better understand machine behavior and to investigate alternative strategies; this trend should continue at NLC.

General areas of application for modeling and simualtion would include:

- Machine commissioning where the objective is to reconcile the model with the accelerator.
- Routine machine setup based on the design models and specifications.
- Model driven feedback system as described above.
- Near-automated diagnostic capabilities such as emittance measurement, and lattice diagnostics to quickly identify the sources of machine performance degradation.
- Model-based optimization tools such as orbit and lattice properties correction applications to allow rapid fixes for performance degradation.
- Creation of “multiknobs” which allow one to vary a parameter which may depend on many hardware values in a linear or nonlinear manner.
- Ability to use the same model base for online as well as offline machine physics studies.

14.2.7 Accelerator/Detector Coupling

The interaction region is precious not only to the physicists taking data with a detector but also to those running the accelerator. Certain diagnostic information can only be obtained there, but instrumentation is likely to impinge on the detector volume. The situation can be ameliorated somewhat by close coordination between shift personnel in the Accelerator Control Room and the active Detector Control Room as well as direct communication between the control system and the detector data acquisition.

At NLC provision should be made for at least the following:

- Shared timing. The Detector Acquisition must be synchronized with the bunch train. It should also have access to the configuration of the train (number of bunches, spacing). In order to make correlations offline each bunch train should have a unique identifier available to both the Control System and the Detector Acquisition.
- Interlocks. Detector components known to be vulnerable must be in a protected state during potentially damaging accelerator tuning.
- Tuning information. The Detector Acquisition should make available to the Control System a suitable collection of background signals at bunch-train rate during periods of good luminosity or fine tuning. This information is needed both for online tuning and for offline analysis.
- Veto information. The Control System should promote knowledge of anomalous pulses to the Detector Acquisition front-end.
- Polarization state. Both the Control System and the Detector Acquisition need access to this.
- Slow updates. Various quantities – state information, statistics – kept on one side may be of interest to the other.

14.2.8 Reliability and Availability

The availability of Control and Instrumentation Systems in existing accelerators (FNAL, SLC, CESR) has achieved the availability numbers required for the NLC (98-99%). The NLC is 10 times larger, and its construction comes 10 years (or more) after the existing machines. From the reliability point of view, the fact that the NLC is larger is the engineering challenge; the fact that its construction occurs a decade later is the engineering opportunity.

The NLC Control System will use components from the computer, communications and electronics industries. These industries have a long-standing record of improving their product reliability and availability in response to market needs. Thus, one can expect that industry will supply improved component reliability which will, in turn, get us part way to the achievement of our availability goals.

It would be wonderful if we could leave it up to the marketplace to solve this problem. Unfortunately, we cannot. We have already seen in existing accelerators that availability is heavily impacted by local design and operating decisions. Spares availability, and the scheduling of time for checkout and maintenance debugging are examples of local operating decisions. The use of redundancy in designs, and the design of equipment in ways to decrease the MTTR (Mean Time To Repair), such as hot swapping of spares, are examples of local design decisions.

We have not yet expended engineering effort in putting together a design and operational strategy to achieve the availability goals of the control system. However, the goals appear to be realistic in terms of the experience of presently existing accelerators. The technologies involved for the control system are well known and tested. There are many practitioners in the fields of reliability, availability, and maintenance of large complicated systems. In the fields of electronics, communications and computers, this is a well-established engineering discipline. We know that effort in this area has to start early in the life of a project, since reliability goals must be set and given to designers early in the design phase.

14.3 Architectural implications

To zeroeth approximation, NLC is just a large accelerator requiring a large control system. To this extent NLC can use something like the control system "standard model" prevalent in large modern control systems. But it is also clear that the functional requirements discussed in the previous section can only be satisfied by making substantial perturbations to this model. These fall into the following categories:

Network bandwidth NLC's heavy dependence on feedback will make significant demands on realtime per-pulse bandwidth between device-controlling computers (called "micros" at SLC). There is also a need for real-time communication in order to synchronize measurements. Finally, to employ online modeling and provide information for simulation, large amounts of data must flow from the micros to an arena accessible to online servers and perhaps to a logger. That is, the control system will have a substantial data acquisition component.

Computational resources As pointed out in Section 14.2.1 a single fast feedback loop at SLC consumes the better part a micro. By contrast, less than 100 micros (mostly running at a lower clock rate) are used for standard device monitoring and control of the thousands of devices comprising the control system. The lesson for NLC, with its anticipated heavy use of feedback, is clear. Online modeling will be another cpu-intensive activity.

Interface with detector data acquisition In order to support the communication outlined above in Section 14.2.7, well-defined interfaces (hardware and software) and adequate network bandwidth between the NLC control system and the detector data acquisition will be needed. This kind of communication has proven essential for SLC/SLD (for example, parts of the detector readout are of great value as accelerator diagnostics). but was absent from the original design. The ad hoc methods currently employed lack flexibility and the amount of information which can be transmitted at 120 hertz is pitifully small.

Hierarchical software organization The cost of the control system for the SLC was much larger in proportion to total project cost than that of any other machine yet built. Already 200 to 300 person-years of software resources have been invested and further improvements are planned. This software development effort has been primarily an intellectual challenge as opposed to a bookkeeping or task management exercise. In many ways, the pace of evolution of the control system has been limited by the time required for the machine physicists to identify and understand a problem and specify a solution.

To appreciate the need for hierarchical organization within the NLC control system it is instructive to look at SLC as it has evolved. SLC's database is in many respects the heart of its control system. During the first years of its existence (1980–1986), approximately 80 kinds of objects (with typically many instances of each type) were defined in it. All but a handful describe hardware. Since that time about another 90 kinds of objects have been added to the database, but of these well over half describe something bearing little resemblance to a physical device, for example feedback loops or model parameters. The database, like the system and application software which access it, is both substantially larger and different in character from initial expectations.

This evolution at SLC was only possible because of the generic nature of the treatment of the lowest level components. Strict rules are followed in all low-level device interfaces, and this structure facilitates continued growth through the application of higher and higher layers. This implies that the design of the lower levels is critical for a control system which must support a complex variety of high- and very high-level applications.

With its tight tolerances which must be actively maintained, and opportunities for subtle interactions among its elements, NLC will continue this trend. In order to get their job done efficiently all who use the control system (operators, machine physicists, hardware maintenance personnel, programmers) must be able to access it at an appropriate level of abstraction so they can control, monitor and analyze the entities of interest to them.

14.4 The Control System Model

As noted in the Introduction, we will describe a control system implementation model that is, we believe, a realistic model for an NLC Control System, if it were being built today. It is realistic because it follows the main thread of control system design used in accelerators recently built (CEBAF, ALS) or under construction (PEP-II). However, we will state the standard disclaimer that this is a design based on current technology, and that the real NLC control system design will be based on the technology available at the time of its construction.

This model is based on what has become a fairly standard approach in modern computer control systems (Figure 14-1); it comprises a set of consoles and servers linked to each other and to a hierarchically lower set of front-end computers via high speed networks. These front-end computers are, in turn linked to devices via dedicated control networks. The devices themselves are expected to often include embedded computers; thus the front-end computers are really communicating with still another architecturally-lower layer of device computers. In the case of low to medium multiplicity devices (1 to 1000 devices), there may be a device interface crate (VME, VXI) which contains a module which controls the device. We expect that the truly high multiplicity devices (Beam Position Monitors, Magnet Movers; greater than 5000 devices) will probably have a cost-optimized design which includes dedicated embedded computers, and communicates to the front-end computers via a digital network link of some kind.

14.4.1 Operator Consoles

Console Computers will be located in the Control Room and will be the Primary Operator Interface (OPI) into the Control System. There will be a Graphical User Interface (GUI) driven by the EPICS Control System based on the X-protocol. The physical hardware will be comprised of a processor device driving (perhaps) two video heads with some pointing capability plus a standard keyboard (Figure 14-2). These workstation processors will have large internal memory and internal hard disks.

These machines are modeled as DEC Alpha Processors running the NT Operating System. These machines are capable of high throughput, excellent number crunching, and a fast network response.

Consoles will use the TCP/IP over Ethernets which are in turn connected to the FDDI backbones via switching bridges. Separate local networks will be used to subdivide the overhead status display screens and the pairs of screens associated with each console such that network or server failures will not bring down the control room facilities.

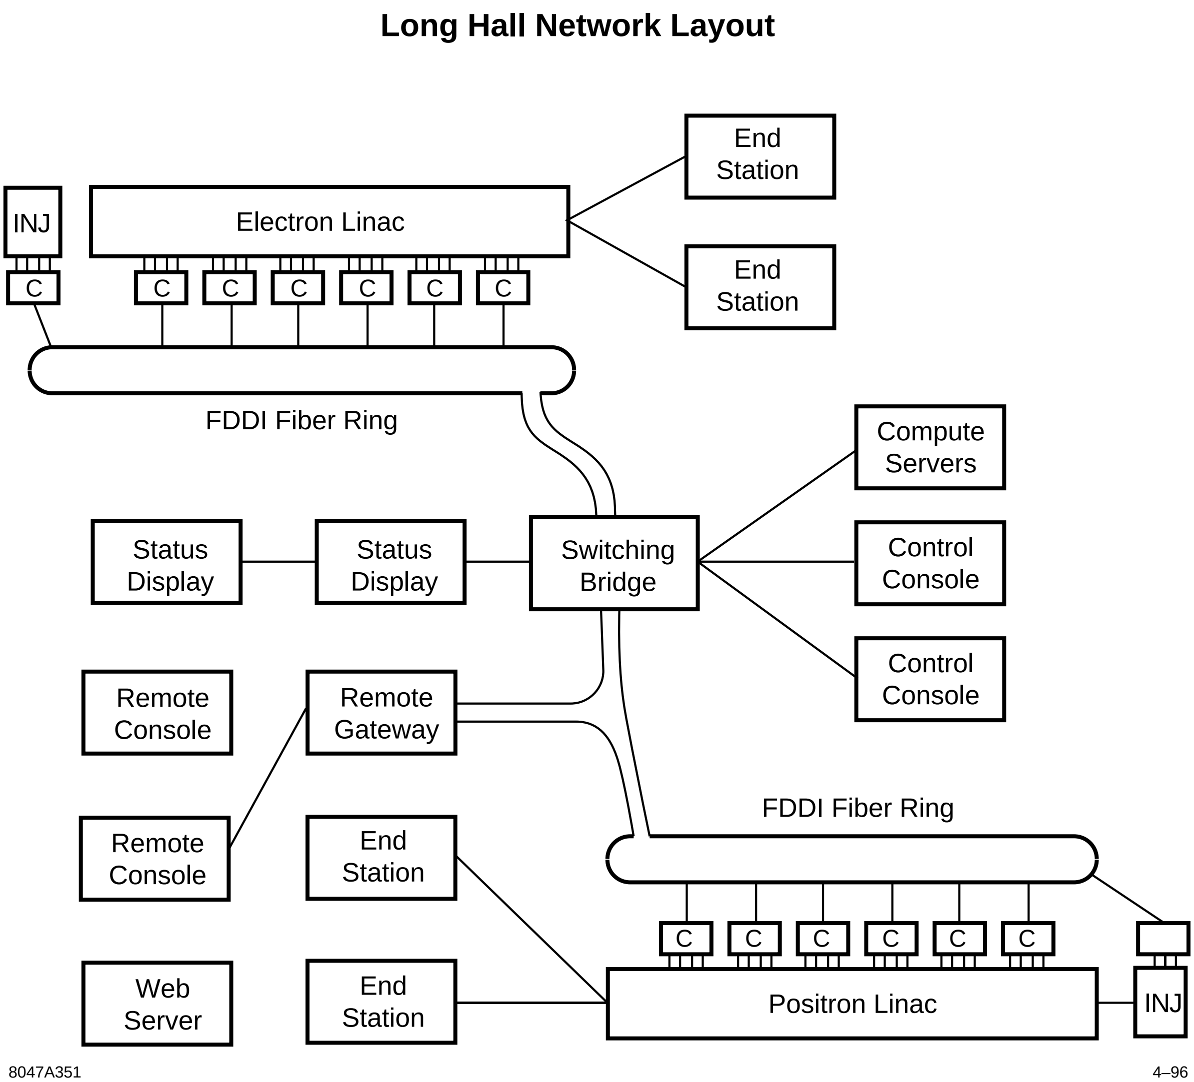

Figure 14-1. *Schematic of the control system layout.*

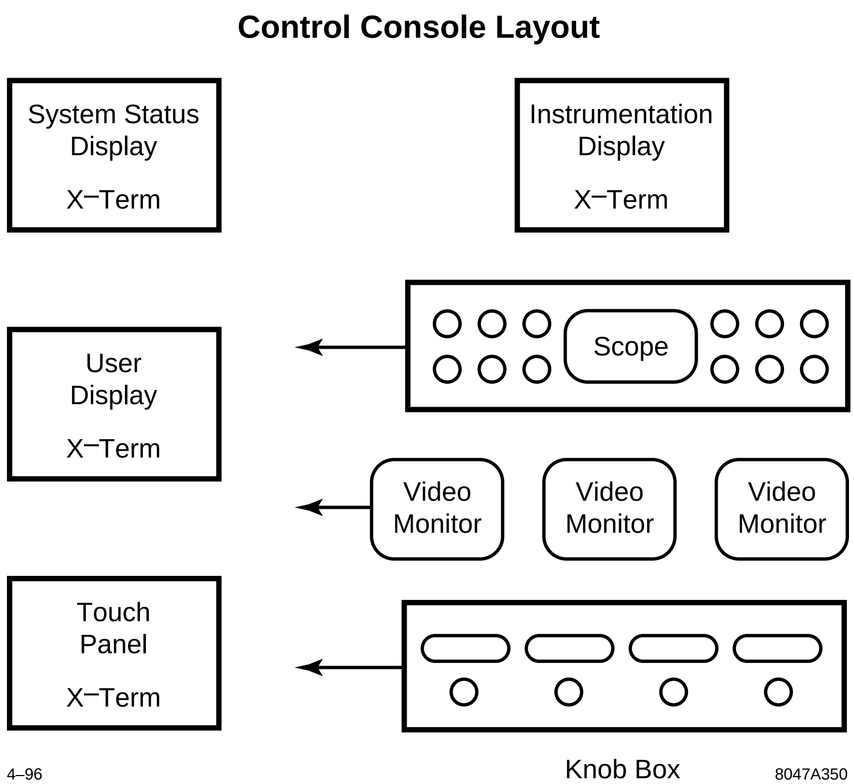

Figure 14-2. *Schematic of operator's console.*

14.4.2 Application and Server Computing Resources

The control system model is a distributed system, and one component of that distributed system will be a set of around 20 computers which will function as servers of various functions. The functions we have in mind are database, network interface, computational, and file servers.

Database servers will include the Control System Database and the Documentation Database (drawings of devices and cable-plant, wire lists, system documentation, operations guidelines, code management, etc.). Network interface servers will handle connections and security issues for computer communication access outside the control system. Computational servers will supply sufficient computing power for accelerator modeling and simulations. File servers will provide storage for a number of functions, including software and firmware development, the archiving of data from the control system of device histories, operation histories, configurations, alarm message and error reporting, computer system and network management monitoring data, etc. Many of these functions will be accessible to the detector data acquisition; conversely, information from the detector will be accessible to control system applications.

Associated with these system will be large disc farms for storage of the large amounts of data associated with such an accelerator complex. Some form of redundancy such as clustering will allow the server complex to gracefully degrade in the event of individual server or other component failure. We expect that this server cluster will be in relatively close proximity to the control room, so that local very high speed links connecting to the consoles in the control room may be utilized.

Network Security Systems

There are requirements to safeguard the Accelerator from remote access and control by unauthorized individuals. Security systems will be implemented at the onset of the program which allow data to be freely available over external networks, while preventing unauthorized remote users from operating equipment. These security systems will provide encrypted remote sessions such that passwords and private data cannot be swept for unauthorized use.

14.4.3 High Speed Networks (FDDI & Ethernet)

This Distributed Control System depends heavily on networks with substantial bandwidth to function properly. Thus it is critical to use network facilities which can handle anticipated loads at low utilization figures. These networks need to be industry standard facilities as well.

Accelerator Control Networks will be subdivided into subnets to isolate the operation of the machine from another networks activities at the laboratory. There will be a control network Firewall router between general laboratory networks and the control networks to allow isolation if required.

In terms of current technology , FDDI and Ethernet make good sense because they are industry standards, they are versatile, and their performance is well understood. At the time of the final design review, it will be necessary to review these specific selections and to use the then current stable dependable technologies.

There are some special considerations for the networks, in that the length of the network runs will generate challenges for any network technology. Fiber implementations will be important, and microwave facilities may be an option to consider.

14.4.4 Front-end Computers

Front-end computers are the Input/Output Controllers (IOC) which actually perform the data acquisition and control in real-time in the accelerator housing. These machines are microprocessor based platform using PCI, and VXI bus backplanes. Each device will have its own Real-time Executive, and will run a standard set of acquisition/control software. Database configurations will be downloaded in order to structure the number and type of devices and processes the controller handles.

These processor crates will normally be operated remotely over the Ethernet from the Control Room, but can be operated locally from a terminal for diagnostic purposes. Long-haul communications will be handled by the FDDI backbones with bridges to distribute Ethernet connections to micro crates, local control X-terms, and other Ethernet devices.

14.4.5 Data Acquisition Crates

Data Acquisition Crates will be a mixture of VME and VXI crates containing the analog and digital input/output cards used to make measurements and define set points. Low-cost VME crates will be used for most of the modules, along with some industrial equipment for the low accuracy systems. VXI will be used for the precision measurements and high-frequency rf and timing modules.

14.4.6 Instrumentation Modules (VXI & GPIB)

There will be a number of modules for high frequency monitoring and control which will emanate from VXI and GPIB controlled instruments. The VXI crates will be controlled over Ethernet through the controller module in the crate. These crates can control locally positioned GPIB instruments as well.

14.4.7 Radio Frequency Control, Phasing & Feedback

The rf requirements are covered in Chapter 8; we will not repeat those here.

14.4.8 Dedicated Control Networks

These networks are dedicated point-to-point networks to connect specific pieces of equipment. This implementation is in use at SLC has been chosen as perhaps the only one available today to get high performance and dedicated functionality.

These links will be found between control crates and large power supplies with self-contained controllers. There will also be a link or links between the control system and the detector data acquisition in order to transmit, for example, tuning information (detector to accelerator) and beam quality information (accelerator to detector). There will be point-to-point links between machine stabilizing feedbacks and feed forwards. This approach (point-to-point links) has serious drawbacks for feedback: it does not scale well and lacks flexibility. Alternatives should be investigated as they become technologically feasible.

Special 1553 links have been employed between elements of the SLC machine protection system for security and speed.

14.4.9 Timing and Beam Rate Control

The purpose of the scheduling system is to provide control of the paths to be followed by beam pulses more or less in real time. Several pulsed beam dumpers and diagnostic stations will be installed throughout the NLC; before and after the damping rings, in the injector and positron system, after the bunch compressor, at several places throughout the linac and on either side of the big bend and collimation systems. The scheduling system will control the firing of these dumper magnets and synchronize the data acquisition on the dumped pulses. It will also be used for the machine protection system (see Chapter 16).

The scheduling system is modeled on the one currently in use at SLC. It will be used to control the injection, extraction and storage time in each of the three rings. Using it, the control system will be able to program the storage time, and therefore the output emittance, of each of the three rings. This feature will be built into the system through a linking, or pointer-based, structure that will be used to track the progress of a bunch or train of bunches from it's inception through to a dump. Another requirement of the scheduling system is to provide control of synchronized data acquisition and sequencing of pulses when a fast pulsed device, such as a kicker magnet or a pulsed phase shifter, is being adjusted. Such acquisitions will be used to quickly scan a beam across a beam size monitor and for certain classes of beam optimization. They are also required for optimizations which involve measuring derivatives such as maximizations or minimizations or maintaining a specified phase with respect to the rf. The technique to be used is a synchronous detection scheme with a sub-tolerance dither.

Accelerator Beam Rate will be controlled by a dedicated special purpose processor running custom software to support flexible rate control in both accelerators. This Master Pattern Generator will be wired into the machine protection systems to handle both rate limiting and machine protection shutdown. Rate information will be disseminated to appropriate components of the detector data acquisition as well as as to front-end computers, etc., within the control system.

Flexible rate control allows controlling the amount of beam energy transported throughout the machine. This will facilitate the alteration of beam parameters for experiments which require unique beam characteristics or timing.

This flexible beam rate control also allows the accelerator to be rate limited by classes of machine protection problems so that problems can be located and identified with low rate beams with reduced risk of accelerator or equipment damage. Beams will automatically rate limit back to designated rates as problems are resolved. Rate control will be exercised on a pulse-to-pulse basis.

14.4.10 Machine Protection Systems

The purpose of the machine protection system (MPS), which is described in Chapter 16, is to prevent damage to machine system components in the event of a routine failure. The system is not intended to provide comprehensive protection against any possible failure. One of its main functions is to automatically provide a sequence of beam pulses that can be used as effective diagnostic tools during a startup or fault period.

The machine protection system has four logical layers: 1) mechanical, 2) device controllers, 3) power monitoring and 4) beam scheduling and control. The layers provide a graded approach that allows the production of a nominal intensity single bunch beam for diagnostic purposes. All mechanical systems should be capable of withstanding the impact of a single pulse of such a beam. They should also be able to survive another strike in the same location. It may not be possible to develop structures that can stand a single pulse strike from a nominal single bunch at nominal emittance, especially for the higher intensity versions of this design. In this case, an emittance enlarging system must be integrated with the MPS so that proper operation is ensured. Once operation is checked with single bunch beams, the repetition rate may be increased, the emittance brought to nominal and the number of bunches brought to its full value. This sequence must be applied in this order. The only viable way to transport full-intensity beams is to make sure that the transverse deflecting forces acting on the full-power beam cannot change enough in the interval between pulses to target the beam cleanly on a beamline element. One consequence of this is that low-repetition rate, full-train intensity operation, is not possible.

14.4.11 Equipment & Tunnel Access Control

The control system will monitor the status of equipment and the state of Accelerator access, but the actual control of Personal Protection, Machine, protection, and Hazards will be handled by dedicated hardware and Programmable Logic Controllers.

14.4.12 Application Software

Application packages will reside and run on the console processors and on a separate applications processor in the control cluster. This extra applications processor will take on large resource computing loads that would not run well on console machines. Other applications or analysis programs will be off-loaded to user machines via self-describing data files.

Applications packages will include Accelerator diagnostic packages, measurement packages for things like emittance and chromaticity, energy management (LEM), power steering with machine optics models, correlation plots, simulations, machine models, multiknob control, and data archiving.

Measurement packages may be operated remotely at lower priority than control room activities. Analysis of data may be run on remote hosts or the application processor in the cluster. Operation of applications which operate accelerator equipment or change machine configurations or settings will execute from the control room only.

The applications environment will be structured to enhance the ability of the Laboratory to use software developed at other Laboratories or purchased commercially.

Included applications: archiving, correlation plots, steering, LEM, models, emittance, logging.

14.4.13 Software Application Bus

Applications which run on the Console or Cluster Computers will run on a software layer which will isolate them from the complexities of where data comes from, how it is stored, and how it is transported. Hidden facilities will deal with correlation plot data which has to be correlated in time and take into account measurements taken in different parts of the accelerator.

This software structure will make available common measurement and data collection facilities which may be required by application or display processes. Data will be presented and exported in self-describing formats compatible with application packages available at other laboratories [Watson 1995].

Similarly, data files will also be available in Matlab format for local or remote analysis.

14.4.14 Software Development Environment

Software development will be accomplished on workstations similar to those used in the production environment, however, they will be configured to run a parallel but separate control environment so that actual production equipment will not be controlled by accident. With the exception of the separate environments, the software environment will be identical to the production environment. Some special hardware will be developed which will help simulate accelerator operation for software evaluation and testing. Additional software will be required to compile code, to control progressive versions of software, and database facilities to build run-time databases and configure equipment.

Diagnostic systems will include remote diagnostic and debug capability for all networked microprocessor systems (not including embedded systems). Isolated or low device count GPIB instruments will be controlled via Ethernet by GPIB network control boxes placed in the locus of the GPIB instrumentation.

14.4.15 Industry Standards

The Control System will utilize as much commercial equipment and software as practicable. Industry standard equipment and facilities will be utilized to reduce cost and improve maintenance and reliability. To the extent possible, the control system will utilize electronic modules available commercially and utilized in other laboratories to reduce cost and resources involved in hardware development and in writing low-level software drivers.

References

[Humphrey 1992] R. Humphrey, "Lessons From the SLC for Future LC Control Systems," *Proc. of the 1992 Int. Conf. on Acc. and Large Exp. Phys. Control Systems*, KEK, Tsukuba, Japan, KEK Proceedings 92-15 (1992).

[Watson 1995] C. Watson *et al.*, "cdev, a Common Device API," *Proc. of the 1995 Int. Conf. on Acc. and Large Exp. Phys. Control Systems*, Fermilab, Batavia, IL (1995).

Contributors

- Joanne Bogart
- Spencer Clark
- Rusty Humphrey
- Nan Phinney
- Marc Ross
- Hamid Shoaee