SSCL-394

SSCL-394

Superconducting Super Collider Laboratory

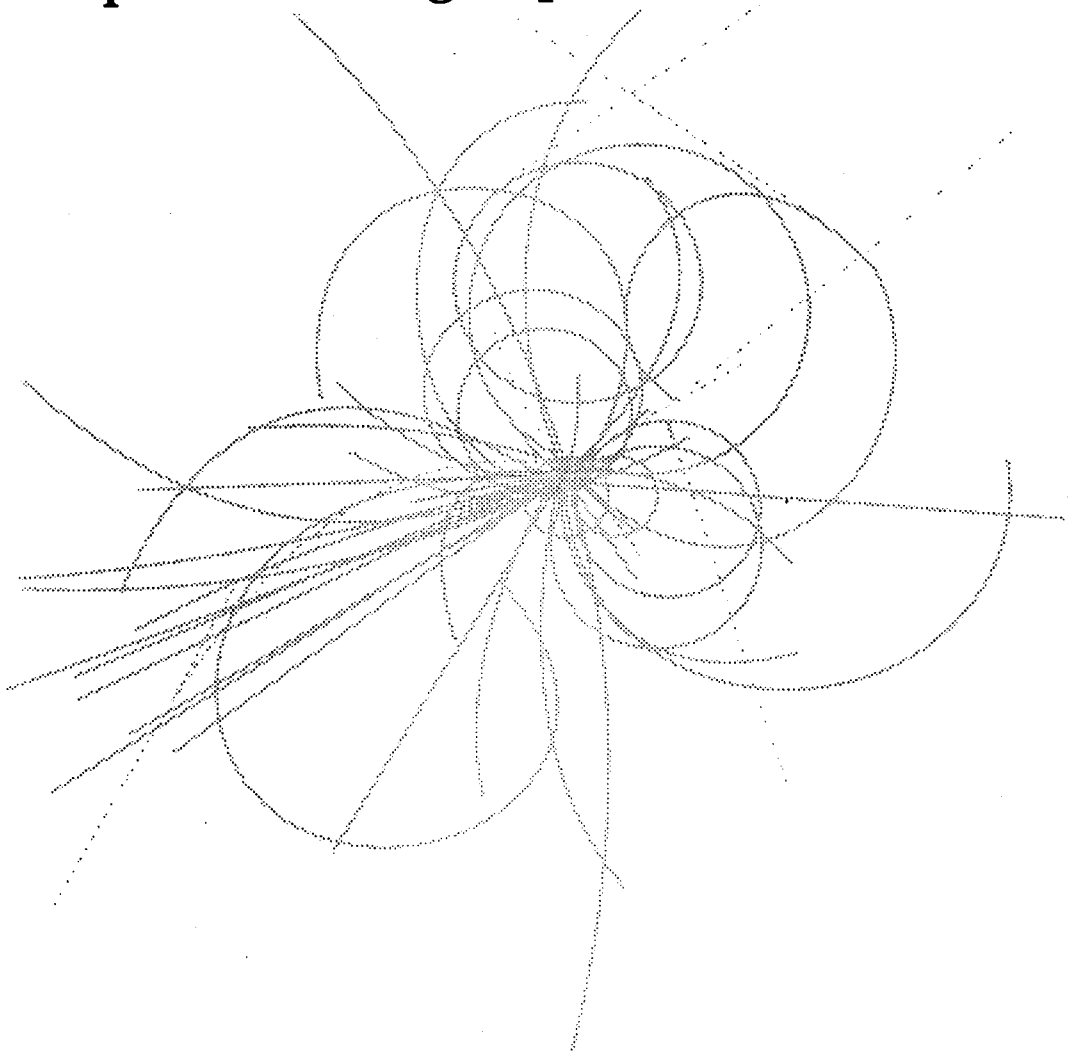

A High Performance Architecture for Accelerator Controls

M. Allen, S. Hunt, H. Lue,
C. Saltmarsh, and C. Parker

March 1991

To be published in Supercollider 3, Plenum Press, New York

SSCL-394

A High Performance Architecture for Accelerator Controls*

M. Allen, S. M. Hunt, H. Lue, and C. G. Saltmarsh

Accelerator Division
Superconducting Super Collider Laboratory†
2550 Beckleymeade Ave.
Dallas, TX 75237

C.R.C.B. Parker

SL Division
CERN
CH1211 Geneva 23
Switzerland

March 1991

---

*Presented at the 1991 International Industrial Symposium on the Super Collider, Atlanta, Georgia, March 13-15, 1991.

†Operated by the Universities Research Association, Inc., for the U.S. Department of Energy under Contract No. DE-AC02-89ER40486.

A HIGH PERFORMANCE ARCHITECTURE FOR ACCELERATOR CONTROLS

M. Allen, S.M. Hunt, H. Lue, and C.G. Saltmarsh

Accelerator Division
Superconducting Super Collider Laboratory*
2550 Beckleymeade Avenue
Dallas, TX 75237

C.R.C.B. Parker

SL Division
CERN
CH1211 Geneva 23
Switzerland

Abstract: The demands placed on the Superconducting Super Collider (SSC) control system due to large distances, high bandwidth and fast response time required for operation will require a fresh approach to the data communications architecture of the accelerator. The prototype design effort aims at providing deterministic communication across the accelerator complex with a response time of <100 ms and total bandwidth of 2 Gbits/sec. It will offer a consistent interface for a large number of equipment types, from vacuum pumps to beam position monitors, providing appropriate communications performance for each equipment type.

It will consist of highly parallel links to all equipments: those with computing resources, non-intelligent direct control interfaces, and data concentrators. This system will give each piece of equipment a dedicated link of fixed bandwidth to the control system. Application programs will have access to all accelerator devices which will be memory mapped into a global virtual addressing scheme.

Links to devices in the same geographical area will be multiplexed using commercial Time Division Multiplexing equipment. Low-level access will use reflective memory techniques, eliminating processing overhead and complexity of traditional data communication protocols. The use of commercial standards and equipment will enable a high performance system to be built at low cost.

*Operated by the Universities Research Association, Inc., for the U.S. Department of Energy under Contract No. DE-AC02-89ER40486.

REQUIREMENTS OF THE SSC LAB CONTROL SYSTEM

The control system of the Superconducting Super Collider Lab (SSCL) will need to operate at high speed over large distances and offer a fast response time. The data acquisition and control requirements have been categorized,[1,2] and a maximum response time of less than 100 ms is required. Data rates will vary from 2 bytes/sec for vacuum equipment to as much as 50 kbytes/sec for some beam monitoring equipment. More than 150,000 control points must be accessed, including those for the ramp generators, superconducting dipoles, superconducting quadrupoles, vacuum pumps, RF systems, cryogenics, and cooling water. Total data acquisition and control bandwidth requirement approaches 2 Gbits/sec.

Existing large accelerator control systems often use local area networks for equipment access. These were designed primarily for peer-to-peer communication, rather than the hierarchal communication that characterizes much accelerator control. In an accelerator there is generally little data exchange between local processors.

Contention-based CSMA/CD local area networks such as Ethernet cannot guarantee any network station access to the network at any fixed or even minimum rate. This type of network is said to be non-deterministic. CSMA/CD networks also exhibit rapid degradation of performance under heavy load, and suffer from a limited geographical coverage because the propagation delay of the network must be shorter than the time of transmission of a minimum-size packet, leading to inefficient transmission of short messages.

Token-passing networks, while providing deterministic communication, have not adequately met the needs of large accelerator control systems due to reduced response time as the number of stations and size of the network increases. Furthermore, many layers of protocol running on non-real time operating systems have led to unpredictable performance as seen by application programs.

The control system requirement to provide 100 ms response time to all equipment—at whatever data transfer rate is required for that equipment—precludes the use of token passing or CSMA/CD local area networks.

TIME DIVISION MULTIPLEXING

The requirement for high-bandwidth, deterministic communication suggests a highly parallel system of data communication operating in the time or frequency domain. Frequency division multiplexing interfaces are not widely available for general purpose computers and micro-controllers. Time division multiplexing equipment, however, is widely available and extensively used in the telecommunication industry.

Time Division Multiplexing (TDM) is used by telephone companies for the efficient transmission of digitized voice and data traffic. It can also be used as a means of sharing a physical medium (fiber or copper) between different forms of data,[3] such as digitized voice, packet switched data, Ethernet, Token passing networks, digitized video, and others.

Standards have long existed for low-speed TDM interfaces:[4] T1 (1.544 Mbits/sec) is the equivalent of 24 voice channels (each 64 kbits/sec); DS3 (45 Mbits/sec) can carry 24 T1 links. Standardization has now been achieved at higher data rates by the Synchronous Optical NETwork (SONET) standards. SONET is a U.S. standard equivalent to the Synchronous Digital Hierarchy (SDH), defined by the CCITT for international (non-U.S. and Japan) equipment. This standard defines a way of interfacing transmission equipment from different vendors—not only the format of the data on the links, but also management protocols for maintaining the network.

SONET multiplexors are now available up to speeds of 2 Gbits/sec, enabling large amounts of data to be economically transferred over single-mode fiber optic cable. This use of TDM over fiber optic cable will significantly cut down on the amount of copper cable that must be installed in the accelerators.

IMPLEMENTATION OF CONTROL LOOPS

Control loops can be located within the front end interface crates using dedicated microprocessors, regional computers located near the front end electronics, or computers in the central control room. Each of these approaches offers advantages and disadvantages.

Local control loops provide the fastest response time, as no data communication is necessary. This would be the preferred solution where safety issues preclude the integrity of the system from relying on data communications equipment.

Regional control loops would be used where a group of devices must be controlled as a whole but without central intervention. Such cases might be for local commissioning of equipment before the central control room is operational.

Global control loops would be provided by centrally situated computers partitioned by function (rf computer, vacuum computer, access control computer, beam monitoring computer). This approach would simplify maintenance, as most complexity could be grouped at a central location. Application programs would operate faster and be simpler to write as they would be operating in their own program space. This approach would be especially efficient where an application required data from more than one location.

DATA TRANSPORT PROTOCOL

The use of high speed TDM links within the accelerator does not by itself ensure a high-performance control system. If many layers of protocol are used between the application and the data transport medium, the control system is slowed both by the delays in implementing that protocol and by the software overhead of processing the vast number of communications links that are needed. What is needed is a way of reducing both the number of protocol layers and the processing overhead of the data communication.

REFLECTIVE MEMORY

Reflective memory is the transferring of data between processors in a way transparent to the application software using that data. Typically a process will act on (read or write to) local memory, the contents of which are copied to the memory of another, physically separate system where other processes can act on it as if it were local to them. Using this technique, processes are unaware of data sharing and do not have to be involved in the complexities of network transmission and protocol. The size of the data unit transferred, speed of data transfer, and update rate can, of course, be varied to suit the applications.

The reflective memory buffer can be shared between two or more processors and can vary in size. The buffer can be sent at a fixed time interval, when requested by a process (typically when an operation on the data is complete), or when any changes have been made to the data. It is also possible to send only the changes that have occurred.

If the whole buffer is sent, then the update rate is low for larger buffers, but the process writing into the buffer can write at full speed. If the process writes to the same location twice before the buffer has been transferred, then only the second value will be sent. However, if only the data written are transferred the size of the buffer does not affect the update rate. But a mechanism must be provided to prevent the process writing data from accessing the buffer at an aggregate rate greater thanthe data transfer rate.

The physical medium for data transfer can be a serial or parallel bus. A parallel bus typically consists of twisted pairs between the crates: 8, 16, or 32 bits wide. The transfer rate is limited by distance. A serial bus has a lower transfer as a parallel bus for the same distance and medium, but can make use of media such as fiber optic cable that can sustain a high data transfer rate over long distances. The use of standard TDM techniques for serial transfer of reflective memory buffers makes the system non-distance limited, able to operate at potentially high data transfer rates and able to share the physical media (often single-mode fiber cable) with other data transfers.

FORMAT OF DATA

Having a large available bandwidth for data communication means that it is no longer necessary to buffer data at the front end acquisition equipment, transferring only a portion of the data that might be of interest. It becomes possible to transfer and store in the Main Control Room all raw data acquired by sensors. If this data is kept in shared memory, an application program requiring readings from a number of sensors does not have to request data from the sensors, await its return, store it in an appropriate format, and then process it. Instead it can simply access the data in its shared memory space, already formatted.

One possible scenario for a data acquisition and control system would be non-intelligent sensors continually acquiring data, perhaps triggered by the message broadcast or precision timing system. The data would be built into a parcel of, for instance, three 16-bit integer values and sent over a 64 kbit TDM link to a central control computer. The values arriving would be placed into shared memory for access by application programs. Depending on the application, new values could be placed in a circular buffer or they could simply overwrite previous values. In some cases, the data might have to be time tagged.

With data arriving in the shared memory space of centralized functional computers from many sensors, from data acquisition crates, and from regional computers, application programs would have access to the memory-mapped control system. That is, an application would behave as if it were directly connected to the hardware.

A mechanism is needed to access and describe the data. One such mechanism being developed by the SSC controls group is the Self-Describing Data Set (SDS) toolset. This is a set of library routines for the creation, access, and manipulation of data sets within file, network, or shared memory segments of a computer system. Names are assigned to collections of data that can then be accessed and manipulated as a whole or individually. A similar concept, Multiple Object Partitioned Structure (MOPS)[5] is employed at CERN for the control of the LEP accelerator. The approach is analogous to the way in which a computer mass storage file system allows access to collections of data by file names, and allows the grouping of collections of files in directories. As in a disk file system, the description of data does not have to be in the same location as the data itself. Furthermore, the data can be updated (perhaps new values arriving from a sensor) without affecting the description of that data. This can be considered as a way of saving the information describing the structure of data that the compiler throws away.

PROGRAM ACCESS TO DATA

In cases where central computers need to share data, either a pre-assigned or on-demand system of networked memory access could be employed. A pre-assigned networked memory access system might use reflective memory to share data. For the short distances involved in the central control room, a commercial parallel transfer reflective memory system might be used.

An on-demand system would allow any computer to have access to the shared memory region of any of its peers. This global networked shared memory could be implemented using standard local area networks and protocols based in much the same way as files are commonly shared across the network now.

IMPLEMENTATION OF THE CONTROL SYSTEM

The elements of the control system (Figure 1) that have been described in this paper are already largely available from industry. These include single mode fiber, SONET multiplexors, and TDM (T1) interfaces for VME computers. TDM interfaces have been built by the control group for the STD bus computers that will act as data concentrators for low-speed control interfaces. SDS libraries are available for describing and accessing data within the control system.

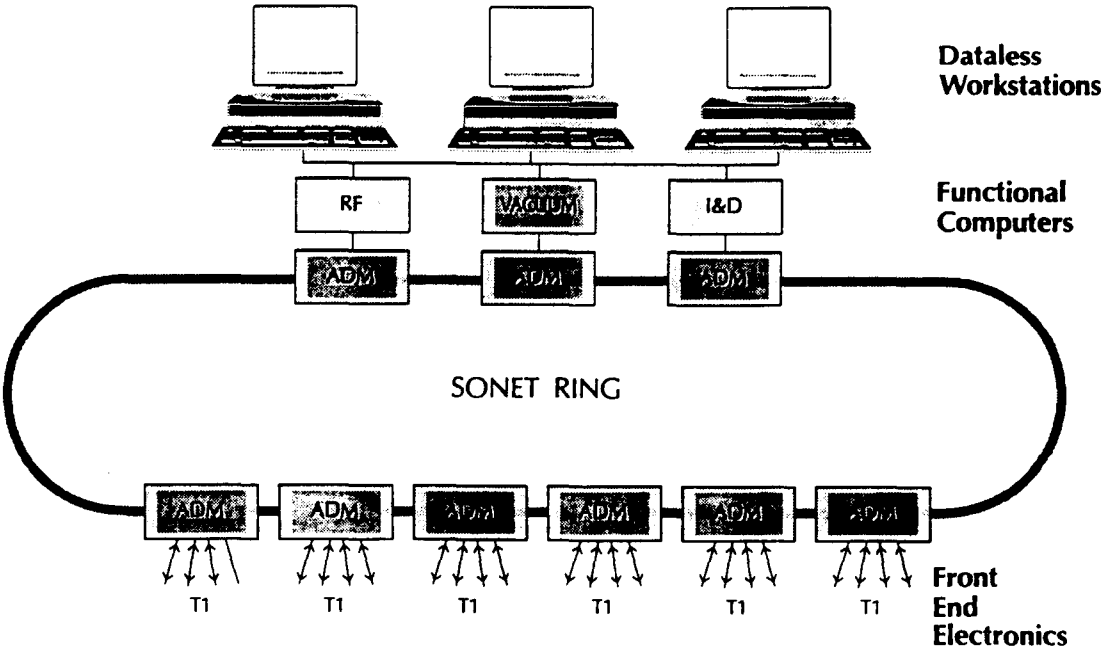

Figure 1. Accelerator Control System Scheme.

Higher rate TDM interfaces, perhaps at SONET rates, will be needed. These should be developed by industry. Standard protocols for program downloading and debugging of real time kernels used in embedded front end applications will have to be ported to run on TDM links. Appropriate error detection and correction protocols will have to be applied to the TDM links.

REFERENCES

1. Site Specific Conceptual Design, SSCL-SR-1056.

2. M.E. Allen, C.G. Saltmarsh, S.G. Peggs, D.A. Dohan, and V.E. Paxson, "Demands Placed on the SSC Laboratory Control System by the Engineering and Operational Requirements of the Accelerators," *International Conference on Accelerator and Large Experimental Physics Controls Systems*, 1989, Vancouver, Canada.

3. M.E. Allen, H.C. Lue, C.R.C.B. Parker, and C.G. Saltmarsh, "A New Data Communications Strategy for Accelerator Control Systems," *International Conference on Accelerator and Large Experimental Physics Controls Systems*, 1989, Vancouver, Canada.

4. ANSI T1.102, "DS1, DS1C, DS2 and DS3 Levels of Digital Hierarchy."

5. W. Herr, "Multiple Object Partitioned Structure," CERN, Geneva CH1211.