'BUBBANET' - A High Performance Network for the SSC Accelerator Control System

S. Hunt, C. Kalbfleisch, K. Low, D. Mathieson
SSC Laboratory *
2550 Beckleymeade Ave.
Dallas, Texas 75237

Abstract

The Superconducting Super Collider Laboratory imposes particularly strict requirements on data networks used in the control and monitoring of accelerator equipment. These requirements are a consequence of the large size (approximately 100 km of accelerators), large number of control points (544,000), and the complexity of the equipment. An overview of the technical systems to be monitored and the projected data rates is presented, emphasizing systems with stringent data communications requirements. We can characterize these requirements in terms of expected network traffic, network throughput or average latency. Analysis of these traffic patterns as applied to different network architectures will aid in identifying the essential components of the final network architecture which meets or exceeds these requirements. We will report on the design decisions and initial results of performance tests on the controls communications network.

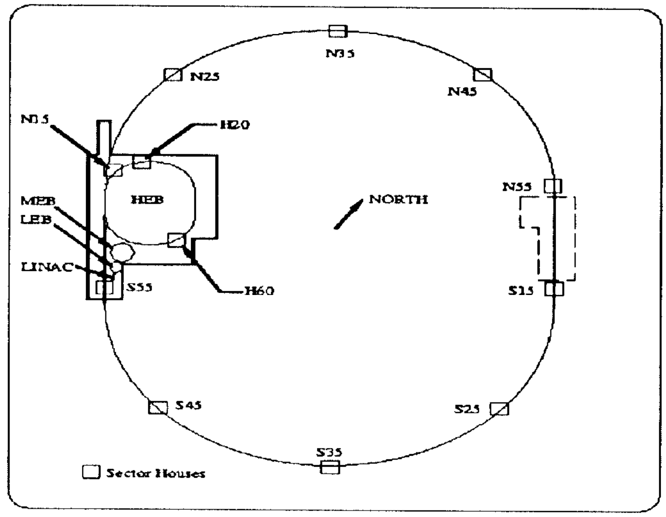

Figure 1: SSCL Accelerators

1 Introduction

BUBBANET (Bidirectional Underground Big Big Accelerator NETwork) is the data communications network to be used to control the accelerators of the Superconducting Super Collider Lab (SSCL). The SSCL consists of a series of proton accelerators being built around the town of Waxahachie Texas. It will consist of a 600 MeV Linear Accelerator (Linac), a 12 GeV synchrotron (LEB), a 200 GeV synchrotron (MEB), a 2 TeV Superconducting Synchrotron (HEB) and a dual counter-rotating 20 TeV Superconducting Collider.

2 Accelerator Control System

The SSC accelerator control system is classified into Beam related controls (RF, Magnet Power supplies, Beam Instrumentation) and Process Controls (Cryogenics, Vacuum, Low Conductivity water). The process controls will not be considered in this document. The SSC Accelerator Beam Controls uses a suite of software tools, originally developed at Los Alamos National Laboratory (LANL), and further enhanced at Argonne National Laboratory (ANS) called EPICS [1]. In addition to LANL and the SSCL, EPICS is also used for controls at Argonne National Laboratory (APS), Lawrence Berkeley Laboratory (ALS), DESY (Tesla) and at Duke University (Free Electron Laser). EPICS has two classes of system, Real-Time VME or VXI systems called Input-Output Controllers (IOCs), and Workstations running UNIX. The SSC accelerator control system is comprised of Front-End Electronics Crates, Supervisory Control Crates and Back-End Crates. The Front-End IOCs provide the first level of control of the technical systems. Supervisory Control IOCs provide real-time control and act as data concentrators for a number of Front-End crates. Back-End systems (Unix Workstations and IOCs) provide high level functions such as operator interface, alarm handling and archiving. The three levels are not, however, enforced. Back-end computers can read data from a Front-End IOC without passing through a Supervisory Crate CPU.

Table 1: SSC Control requirements

Machine	IOCs	Control Points	Size	Cycle Time
LINAC	54	6,000		.1s
LEB	106	23,000	500m	.1s
MEB	120	27,000	4km	3s
HEB	250	58,000	11km	12s
COLL	1450	430,000	87km	3000s

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

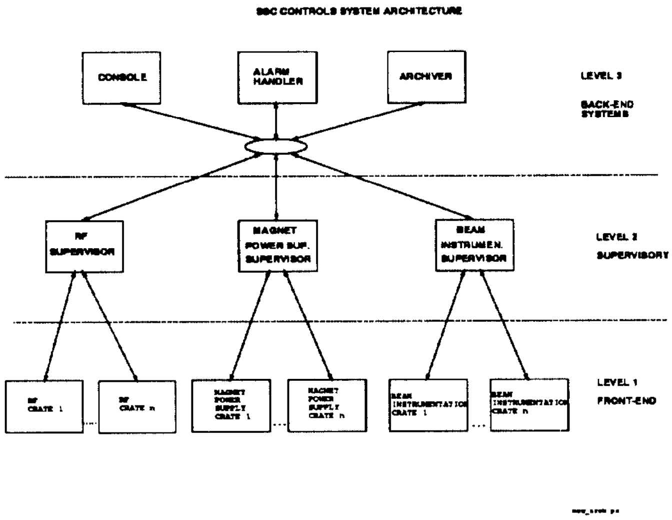

Figure 2: Control system data flows

3 Communications Requirements

The control system has half a million control points and two thousand computer nodes spread around one hundred kilometers of network. The availability requirement for the communication system is less than 10 hours unscheduled down time per year.

Accelerators are characterized by having much of the demand for network bandwidth being concurrent, at times such as injection. At other times, such as during coast, the networks are relatively quiet.

Systems for power control, including Kicker magnets, Corrector Magnets, Ring Magnets and Quench Protection, will have a supervisor crate monitoring a number of slave VME crates at data rates of up to 64 kbits/sec. Ring Magnets and Quench Protection also require deterministic communication with 1ms maximum response time over distances of up to 40km.

RF control crates will be grouped for each machine at a single location. At that location a supervisor will control the groups of three crates required to control each cavity. Data rates are 1 M bit/sec per cavity to the RF supervisor for that machine.

Beam Instrumentation will also have a supervisor, acting mainly as a data concentrator, to avoid each user simultaneously asking for data from the many crates supplying beam information data. Each of the Beam Instrumentation crates requires up to 64 kbit/sec bandwidth.

Total network traffic is greater than 1 Gbit/sec (peak). Much of the data-flow is hierarchical, with little data flow between a front end computer node and its peers.

Because of these factors it was required to design a network infrastructure which provided high reliability, long lifetime and deterministic performance, used commercial components, and had bandwidth management, performance monitoring, fault analysis, and extensibility.

4 Local Area Networks

Commonly used Local Area Networks used to link computers do not meet the requirements of the SSCL. CSMA networks like Ethernet are not deterministic, and are limited in distance and bandwidth. Higher speed token passing networks, such as FDDI, are limited by the single token causing deteriorating response times as network size and number of nodes increases. When sending short packets over large FDDI rings, as is often the case in a control system, actual realisable bandwidth is much less than the theoretical maximum, down to 1% when transmitting 60 byte packets over a 100km ring.

Newer networking technology, such as Asynchronous Transfer Mode (ATM) and Scalable Coherent Interface (SCI), may be appropriate in the future, but are not yet mature, or even stable.

5 BUBBANET Infrastructure

The BUBBANET infrastructure is a layered communication system designed to meet the present and evolving needs of the SSC accelerator controls network users. It is designed on a number of layers designed to be complementary but allowing future requirements to change one layer without modifying the others.

5.1 Fiber optic cable

The first layer in the BUBBANET will be an installed infrastructure of single-mode fiber optic cable. Single-mode fiber was chosen for a number of reasons. Using present day technology, it has a much higher information carrying capacity than multi-mode fiber or coaxial cable. Its latent information carrying capacity is very high, allowing future upgrading of electronics to provide more bandwidth without installing more fiber. Modern single-mode fiber is more resistant to damage by ionizing radiation than multi-mode fiber. Fiber optic cable takes up much less volume than coaxial cable. Single-mode fiber is cheaper than multimode or coaxial cable.

5.2 SONET

Synchronous Optical NETwork (SONET), has been chosen as the transport service running over the single-mode fiber. It is a standard for the transport of Time Division Multiplexing (TDM) channels over high bandwidth links. SONET multiplexors are available from many vendors, at data rates from OC1 (50 Mbits/sec) to OC48 (2 Gbits/sec). SONET has a number of advantages over older generations of multiplexors. It is an International Standard, compatible (at 150 Mbits/sec and above) with Europe. It does not have a practical bandwidth limit; new technology can increase the maximum bandwidth of equipment while still retaining compatibility with the standard. It has a very rich set of built-in network management facilities for monitoring and provisioning (setting up of the network). Being synchronous, low speed data streams can be extracted without disturbing other traffic. A multiplexor providing this 'in line' de-multiplexing is called an add-drop multiplexor (ADM). Slow speed channels (time-slots) can be rearranged in the high speed data-stream, allowing grouping or redirecting of individual channels without switching. This is termed grooming.

5.3 T1

SONET links will link all equipment locations, but the control system interface to the VME equipment crates will not be at SONET rates. The SONET network infrastructure will carry T1 TDM channels operating at 1.544 Mbits/sec. Interface to all VME equipment crates will be at this rate. Each of these T1 interfaces can itself transport 24 of the 64 Kbit/sec point-to-point channels [2]. Each of these channels can have, by using SONET grooming, a different destination. As each of these channels uses a dedicated time-slot on the TDM data-stream, it provides a deterministic response time required by systems such as magnet power supply controllers. This also suits the classical hierarchical controls that characterize many of the SSC systems.

5.4 Higher Level Protocols

The T1 byte-stream is formatted into HDLC frames on each Channel, thus adding the required structure to format 'messages' and providing a checksum. Although point-to-point transport meets many requirements, including low response time dedicated links over the entire geographic area of the complex, some systems require the ability to select, at run time, data from any of the two thousand crates in the control system. To meet this requirement, the links use the Internet Point to Point Protocol (PPP) which is used to set up a TCP/IP service. Thus, with the addition of nodes acting as routers, full connectivity is obtained, and the standard EPICS communication paradigm (which uses TCP) is supported. Commercial routers are available from industry which support PPP protocol over channelized T1 which are compatible with BUBBANET.

6 Performance Testing

A prototype BUBBANET communication system has been built, consisting of SONET Add-Drop Multiplexors, Commercial VME channelized T1 interfaces, and drivers for VxWorks, Lynx and SunOS. Performance has been measured and compared against SSC requirements. Operating under the VxWorks real-time operating system, BUBBANET achieves an application to application response time of 1ms. and throughput per node approaches 200 K bytes/sec with no degradation under maximum load conditions. This meets all the requirements for the SSC control system. BUBBANET has also been successfully tested with the EPICS software that will be used for the control system. The first Accelerator to use BUBBANET will be the Linac, to be installed in July.

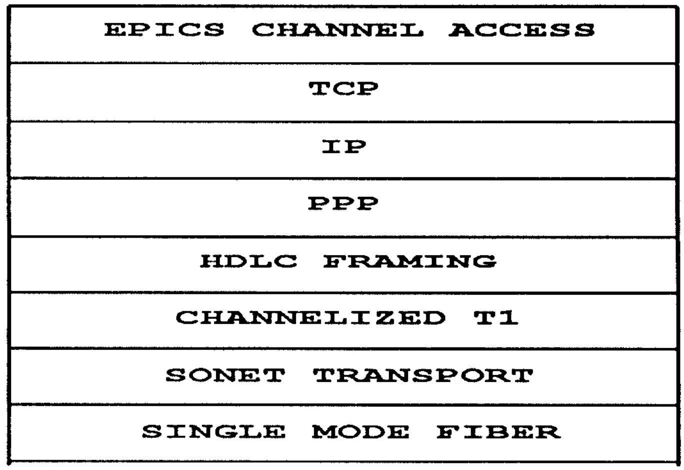

Figure 3: BUBBANET Layers

7 Other Services

The infrastructure of BUBBANET is not restricted to Controls Data Communication. T1 links will also be used to transport live video, and a Message Broadcast System. Live video will use JPEG data compression, to reduce the bandwidth of the video signal for transport over a 1.544 M bit/sec T1 link. Live video will be used in the Personnel Access Safety System for surveillance and for Beam Instrumentation for the transmission of the image of phosphor screens in the beam pipe. The Message Broadcast System [3] will be used to transport timing and synchronization messages around the accelerator complex.

8 Conclusions

The use of international standards and commercial equipment on a layered communications infrastructure, has enabled BUBBANET to be rapidly developed. BUBBANET meets the very demanding requirements of the SSC for reliability, throughput, and response time. The use of SONET has allowed other services such as live video and timing to share the infrastructure.

References

[1] L.R. Dalesio, M. R. Kraimer, A. J. Kozubal, *EPICS Architecture*, International Conference on Accelerator and Large Experimental Physics Control Systems, Tsukuba Japan, 1991.

[2] D. Mathieson, S. Hunt, C. Kalbfleisch, K. Low, *High Speed Serial Communications for Control Systems* , these proceedings.

[3] K. Low, R. Skegg, *Prototype Message Broadcast System for the Superconducting Super Collider*, PAC91.