High Speed Serial Communications for Control Systems

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

Abstract

The Superconducting Super Collider Laboratory is a complex of accelerators being built in Ellis County, Texas. The SSCL control system consists of front-end processors and their associated control points remotely distributed from the Central and Regional control rooms. Control messages passing between these locations require timely (deterministic) distribution. A prototype network consisting of point-to-point links utilizing commercial T1 (1.544 Mb/s) communication boards has been implemented. These dedicated communication links will replace networking services traditionally provided for by shared medium networks like Ethernet(IEEE 802.3) and FDDI(IEEE 802.5). A seamless migration will be achieved by using packet encapsulation based on PPP(Point-to-Point Protocol, RFC 1171). All other networking functions including routing and reliable delivery are still being handled by the usual internet services. A distributed control system that currently uses Ethernet for communication is being re-implemented using these point-to-point links. We will report on throughput measurements, timing constraints and ease of transition to a point-to-point network.

I. Introduction

The Superconducting Super Collider imposes particularly strict requirements for the data networks used in the control and monitoring of accelerator equipment. These requirements are due to the large size of the site and the large number of control points and the complexity of the equipment being controlled [1].

II. SONET and T1

A high speed optical telecommunications technology called SONET [2] will be used to carry data between control computers. The SONET (Synchronous Optical NETwork) standard [3] was developed by the ANSI accredited T1X1 committee, and provides cost-effective, flexible networking using the principles of direct synchronous multiplexing. The use of synchronous multiplexing allows multiplexing of low-speed signals directly onto a high speed backbone without the need for intermediate multiplexor stages. This functionality is integrated with digital switching technology in the Add-Drop Multiplexor (ADM), the device that will be used extensively to connect systems directly to the optical backbone via lower speed T1 (1.54Mbit) links. The network management software built into the ADM allows the creation of logical point-to-point links between any two systems throughout the entire network. In addition to this flexibility SONET is also extremely reliable with every ADM having 100% redundancy including power-supplies and optics, which is important when one considers the availability requirements of the network.

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

III. The Point-to-Point Protocol

The SONET backbone provides the basic point-to-point connectivity between systems, but in order that application software may use this network a networking protocol has to be used. We have selected the Internet Point-to-Point Protocol. The Point-to-Point Protocol(PPP) [4] provides a method for transmitting datagrams over serial point-to-point links. PPP is comprised of three main components:

A. Encapsulation

PPP uses HDLC [5] as a basis for the encapsulation.

B. Link Control Protocol

In order to establish communications over a point-to-point link, each end of the PPP link must first send Link Control Protocol (LCP) packets to configure, test the data link, and negotiate optional facilities.

C. Network Control Protocol

A family of Network Control Protocols [6] (NCPs) are used in establishing and configuring different network-layer protocols. After the link has been established, each host must then send NCP packets to choose and configure one or more network-layer protocols. Once each of the chosen network-layer protocols has been configured, datagrams from each network-layer protocol can be sent over the link.

IV. Implementation

A. Hardware

The hardware currently in use is a commercial VME ISDN card manufactured by Rockwell-CMC that contains intelligent ISDN hardware (Brooktree 8071) capable of generating HDLC formatted frames on any combination of DS0 time slots. In the future we will use a T1 interface built onto a daughter-board conforming to the Greensprings Industry Pak (IP) specification. Several VME-based CPU manufacturers have provided interfaces for Industry Pak's, notably Motorola with their MVME162 which is capable of carrying up to 4 Industry Pak's, each capable of providing a separate T1.

B. Software

In the current implementation on the PPP used at SSCL all of the LCP and NCP have been omitted, since there is no need for the initial configuration and testing of links that will be established indefinitely. During the Boot phase of a system there has been some interest in using LCP in place of a 'bootp'-like protocol when discovering initial configuration options such as the IP-address, since LCP has the ability to propagate this information to remote hosts. Additionally since we need to communicate with commercial PPP implementations within routers, there is a possibility we may need to adhere more strictly to the standard. Device drivers have been written for SunOS (although not using STREAMS for portability reasons), LynxOS and VxWorks [7]. The implementation for VxWorks is to support EPICS (Experimental Physics Industrial Control System) which will be used for the control of all of the technical systems associated with the accelerator complex.

V. EPICS

EPICS has influenced many of the design decisions made during the development of the device driver. EPICS assumes an Ethernet-like network architecture with all hosts able to communicate with any other host. More importantly EPICS assumes the existence of a broadcast facility over all the nodes in the system. It was for this reason that additional functionality had to be built into the driver to allow broadcast messages to be forwarded between nodes. The prospect of forwarding broadcast messages would cause network designers to break out in a cold sweat. We hope that the ability to forward broadcasts will only be used as an interim solution, since the danger of 'broadcast-storms' on a network with as many as 2500 nodes is very real.

VI. T1 Communication

The software has been designed in such a way as to allow a single host to use any fraction of its available bandwidth to communicate with another host.

As can be seen from Diagram 1, the T1 frame from each of the four hosts, A thru D, are split into several channels. Within the ADM, each channel can be routed to either another host connected to the same ADM or onto the optical link, from where it will be removed by another ADM. If a host wishes to communicate with another host which does not have a direct point-to-point link via an ADM, the IP forwarding/routing mechanism will pass the message on to the recipient host via a host that has a point-to-point link both the sender and the recipient. Obviously this form of 'relayed' communication is slow and non-deterministic, and thus only non-critical data would take this path.

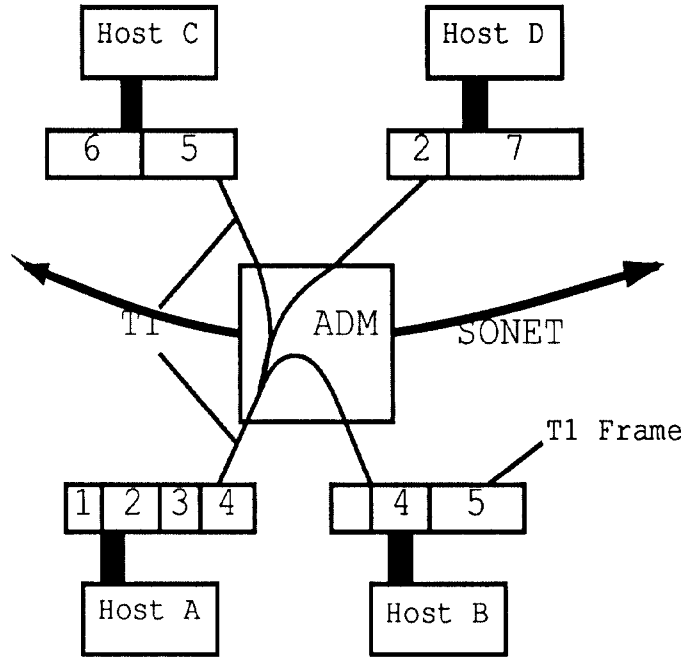

Diagram 1

VII. T1 ⇒ LAN IP Router

It is a requirement that every host be able to communicate with any other on the network. In order to make this possible commercial routers will be used. Most hosts will dedicate a fraction of their bandwidth to a direct communication channel with a router, and, in turn, all routers will be connected together using a high speed backbone (possibly FDDI). In this way any host can talk to any other host with the minimum number of 'hops'.

VIII. Network Management

A. SONET Network Management

Nearly 5% of the SONET data stream is allocated to providing advanced network management and maintenance.

B. SNMP

The Simple Network Management Protocol (SNMP), is a standard method of managing networks from a single console. SNMP allows the remote manipulation of MIBs (Management Information Bases) which hold information about each node on a network. These MIBs hold information about network interfaces, CPU idle time, operating system, etc. In addition, the network administrator may modify values within the MIB to reset statistics or mark interfaces up or down. The description of these MIBs is given in Abstract Syntax Notation 1 (ASN.1) and we are developing a MIB to monitor specific aspects of a SONET network, not covered by the standard MIBs, such as line bit-error rates [8].

IX. Performance

Two areas of network performance are of highest concern to SSC: determinism and throughput.

A. Determinism

One of the reasons for choosing a point-to-point network over more traditional networking technology is the need for deterministic network performance. Shared media networks, such as FDDI and Ethernet, although common in other research centers, fail to provide reliable throughput. In accelerator control systems much of the traffic is 'bursty' with many systems triggered on the same 'event'. This type of traffic is not handled well by shared media networks where all systems compete for the same resource, leading to catastrophic failure unrelated systems due to network congestion. With a point-to-point network, all systems have a pre-determined available bandwidth, which, although usually less than the peak bandwidth available on a shared media network, is completely unaffected by other network users. We have conducted several tests to illustrate the improved determinism of a point-to-point network (1.54Mbit T1) over a more traditional Ethernet. In our test we used four identical systems. Three of the systems, the 'slaves', received a trigger from the fourth, the 'master', after which they sent a data packet to the master using the TCP/IP protocol, as used by EPICS. On Ethernet, normal Ethernet encapsulation was used, and on the T1 network the data was carried using PPP.

Table 1: Performance Measurements

No. Hosts	σ Message Delay (μs)	
	T1	Ethernet
1	220.86	168.26
2	220.20	459.96
3	221.95	777.80

As one can see from the table above, the T1 network performance was unaffected by the number of systems using the network, and the packet arrival time was well bounded. The Ethernet performance, however, demonstrated the contention problems of a shared media network.

B. Throughput

The measurable network performance is unfortunately adversely affected by the current implementation of the T1 interface in the form of a separate intelligent VME card. Taking this into account, we have been encouraged by the performance of the current implementation. We have observed 85% of the maximum theoretical throughput with very little load on the host CPU. These figures are likely to improve with the introduction of a daughter board T1 implementation since much of the buffer manipulation, which seems to be the bottle-neck with the current version, will be more more efficient with a local bus. In addition, the external T1 board has no DMA capability, and the CPU, a Motorola 68020, is much less powerful than that available on the Motorola MVME162.

X. Conclusions

As future network protocols and traffic types emerge, the Point-to-Point Protocol should be able to carry them. If not, the HDLC framing used by the PPP is the same as that used by emerging standards such as ATM, which means that the network infrastructure will not need to be changed to carry other HDLC traffic as well. The use of high speed point-to-point networking is a departure from more traditional accelerator control systems. This decision has enabled us to design a flexible, cost-effective networking solution that will meet the demanding requirement of such a large complex. In addition, the extensibility of SONET technology, coupled with the long lifetime of telecommunications hardware, means that the network will be able to grow to meet any future requirements of the control system.

References

[1] S. Hunt, C. Kalbfleisch, K. Low, D. Mathieson, *"BUBBANET" - A High Performance Network for the SSC Accelerator Control System*, these proceedings.

[2] American National Standard for Telecommunications - Digital Hierarchy - Optical Interface Specification (Single-Mode), ANSI T1.106-1988.

[3] American National Standard for Telecommunications - Digital Hierarchy - Optical Interface Rates and Formats Specification, ANSI T1.105-1988, ANSI T1.105-1990.

[4] Simpson, W. A., *The Point-to-Point Protocol*, RFC 1331, May 1992.

[5] International Organization For Standardization, ISO Standard 3309-1979, *Data communication - High-level data link control procedures - Frame structure*, 1979.

[6] McGregor, G., *The PPP Internet Protocol Control Protocol (IPCP)*, RFC 1332, May 1992

[7] Wind River Systems, Inc., 1010 Atlantic Ave., Alameda, CA 94501.

[8] C. Kalbfleisch, S.Hunt, K. Low, D. Mathieson, *Network Management of Real-Time Embedded Processors*, these proceedings.