An Accelerator Controls Network Designed for Reliability and Flexibility[*]

William P. McDowell and Kenneth V. Sidorowicz
Advanced Photon Source, Argonne National Laboratory, 9700 Cass Ave, Argonne, IL USA

Abstract

The APS accelerator control system is a typical modern system based on the standard control system model, which consists of operator interfaces to a network and computer-controlled interfaces to hardware. The network provides a generalized communication path between the host computers, operator workstations, input/output crates, and other hardware that comprise the control system. The network is an integral part of all modern control systems and network performance will determine many characteristics of a control system. This paper describes the methods used to provide redundancy for various network system components as well as methods used to provide comprehensive monitoring of this network. The effect of archiving tens of thousands of data points on a regular basis and the effect on the controls network will be discussed. Metrics are provided on the performance of the system under various conditions.

1 Introduction

The APS accelerator control system has been implemented using the Experimental Physics and Industrial Control System (EPICS) software tool kit. At the APS, the operator interface is a SUN Microsystems UNIX workstation with an X-windows graphical user interface. As is common to this type of system, an operator workstation or X-terminal may be located at any physical location in the facility. An operator has the ability to generate and alter control displays and to access applications, interactive control programs, custom code, and other tools. The front-end computer or input/output controller (IOC) provides direct control and input/output interfaces for each accelerator subsystem. At APS the standard crate uses the VME or VXI bus standard, a Motorola 68040/60 processor, Ethernet-based network communications, and a variety of signal and sub-network interfaces. The 68040/60 processor provides the crate with the intelligence to allow it to run its software autonomously with respect to all other devices in the system. The EPICS core software running in the crate hides hardware dependencies from the high-level software running on the workstation. There are approximately 175 crates used in the accelerator control system. A real-time operating system, VxWorks, is run in the crate central processing unit (CPU) to provide the basis for the real-time control.

EPICS uses the TCP/IP networking protocol, a commercial standard supported by all network hardware vendors. The TCP/IP implementation is independent of the particular network medium selected to implement the network. APS uses FDDI, 10-Mbit Ethernet and 100-Mb Ethernet.

The APS user community has very high expectations for the guaranteed delivery of beam. The goal for APS operations is to provide greater than 90% availability (actual hours/scheduled hours). Four percent of the scheduled time is required to fill the storage ring, which leaves only 6% of the availability budget for all accelerator systems failures. The downtime budget for the control system is less than 0.05%. During the operational period from 8/19/97 to 9/15/97 beam was delivered for 495 of 560 scheduled hours, for an overall availability of 88.3%. The error budget for the control system would have been 2.8 hours. In fact, the control system accumulated 0.0 hours of downtime during this period. One of the factors behind the high reliability of the control system has been the excellent reliability of the network.

2 Network overview

Figure 1 gives an overview of the complete APS computer network including that portion of the network that is used to control the accelerators. The accelerator control network uses optical fiber to connect satellite network hubs to a collapsed backbone FDDI concentrator system. All the hubs are dual attached to the concentrator using a star connection configuration. The APS network must serve several diverse functions including accelerator control, beamline control, experimental data acquisition, and normal day-to-day computerized office functions such as word processing. A router is used to isolate the APS network functions. The network lends itself to being divided along geographical lines, and thus the network is divided into the control system, the CAT beamlines and laboratory office modules (LOMs), the central laboratory office building (CLO), and the Argonne Guest House (AGH).

The network hub equipment and cabling plant allows an upgrade path to fast Ethernet, Gigabit Ethernet, or ATM technology if future needs show that an upgrade is needed. In addition, all connections and equipment will allow fail-over to redundant paths and equipment.

The same equipment and strategy has been followed for the network hub equipment in the CLO, the CAT areas, and the AGH buildings. Thus the computers installed in offices, labs, and residence rooms will use Category 5 wiring at 10 or 100 Mbit/sec Ethernet rates, while the network equipment will be able to be upgraded to ATM or

[*] Work supported by U.S. Department of Energy, Office of Basic Energy Sciences under Contract No. W-31-109-ENG-38.

Gigabit Ethernet on an incremental basis.

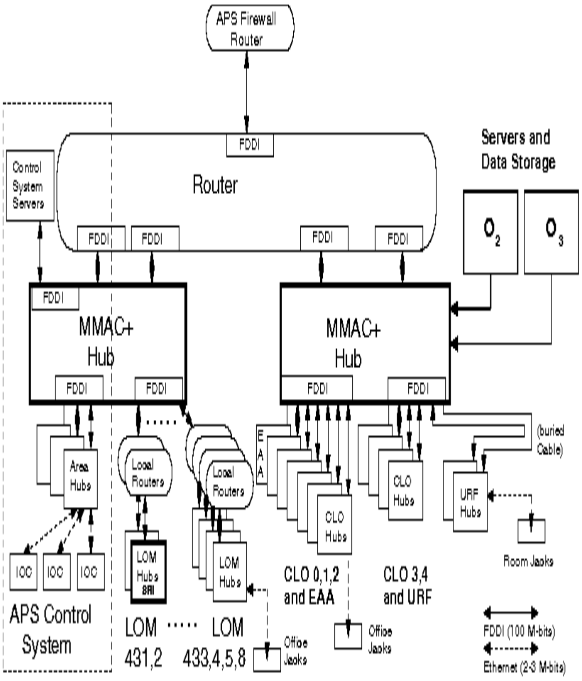

Figure 1. APS network.

3 The accelerator control network

A block diagram of the APS control system hub network is given in Figure 2. An FDDI output card in the APS router is attached to one of the collapsed backbone FDDI buses on the Cabletron MMAC Plus enterprise hub. This hub was selected to provide high availability network services to the APS control system. The MMAC Plus hub accommodates 14 interface modules and has fault tolerant features built into it. It has two dual FDDI networks which provide up to 400 Mbps of network bandwidth, and it will support both packet and ATM cell transport.

There are ten remote hubs in the controls network distributed throughout the accelerator facility in order to provide local Ethernet connections to all network devices. These remote hubs are dual attached to the central MMAC+ hub using FDDI in a star connection configuration. This allows network reconfiguration without the addition of a new fiber plant if future technology, such as Gigabit Ethernet or ATM, is installed. The system also uses different physical paths for the fibers between the remote hubs and the concentrators in order to provide protection against common mode physical damage. There are four hubs serving the storage ring, two serving the rf system, two serving the injector system, and two serving the systems in the main control room (MCR). Each remote hub is attached to both concentrators and each concentrator is completely independent. This independence was one of the primary reasons a Cabletron MMAC+ hub was selected for this application. There is no common management module in the MMAC+; each module in the hub is completely independent sharing only a passive backplane.

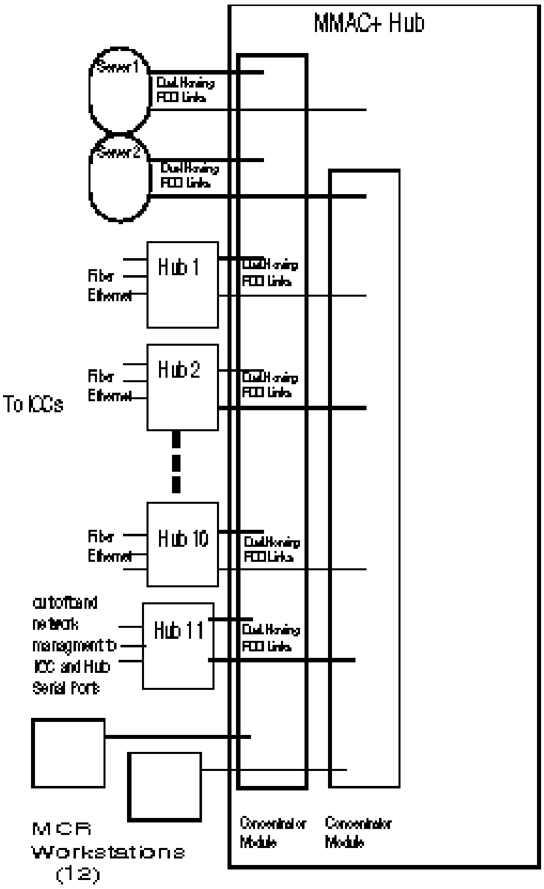

Figure 2. APS controls network.

In addition, two independent, redundant power supplies power the MMAC+ hub. Thus it would take multiple independent failures to cut communication with the remote hubs. All of the control system IOCs are connected to the hubs using fiber Ethernet and they too can be reconfigured if required. To provide redundant service, every IOC is connected to two hubs using a redundant channel fiber optic Ethernet transceiver. Each transceiver is connected to two hubs: a primary hub with fibers running clockwise around the storage ring from the IOC to the hub and a secondary hub with fibers running counterclockwise around the storage ring from the IOC to the hub. This transceiver has two ways of determining that the network has changed. The first is to monitor the link integrity pulse. If the transceiver senses that the link has been lost, it will automatically switch over to its backup port and reflect that change in the LEDs located on the front of the unit. The other method to determine network status is to sense the presence of data on one port of the fiber optic ports. The first port to sense data will be designated as the active port. The transceiver can be configured for either a 2-sec or a on 10-sec quiet period. If the transceiver does not detect data the primary port for the selected time period, the unit will automatically switch to its backup port. When the switch is made, LEDs on the front panel indicate the backup port is now the primary port. This allows a hub to be serviced or to fail without causing the IOC to lose communication with the network. A block diagram of a typical remote hub is shown in Figure 3.

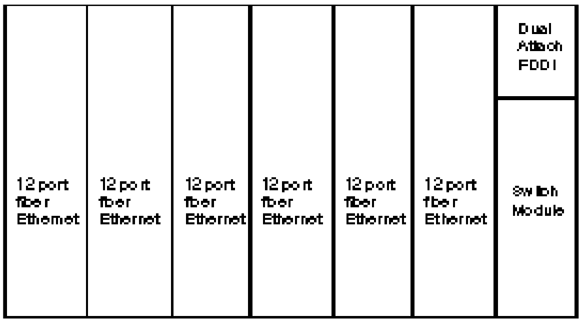

Figure 3. Typical remote hub.

A typical remote hub contains six 12 port fiber Ethernet modules, each of which is connected to an Ethernet switch that in turn connects to the concentrators via FDDI. Each hub module supports six primary IOC ports and six secondary IOC ports. Thus under normal conditions only six IOCs share an Ethernet segment.

A sniffer plot of a heavily used IOC (the main BPM IOC) is given in Figure 4. To obtain these plots APS has installed a distributed sniffer system that allows any Ethernet segment (six per hub) or an FDDI segment to be remotely selected and monitored by a hardware sniffer. The sniffer software supports the X-windows protocol and therefore can be viewed on any X-enabled workstation or PC. Support staff can monitor and diagnose network problems from their office or their home PC using the APS-provided ISDN network link. This allows the system to be diagnosed and repaired without requiring the support staff personnel to travel to the Laboratory during off-hours.

There are twelve workstations in the main control room: seven of these are used by operations to control various aspects of the facility, and five of the workstations are available for general use by engineering groups. Each group of six workstations is single-attached by FDDI to each concentrator so that a concentrator failure will not impact operations.

There are also two file servers in the control system. These computers are dual-attached to each concentrator. Six workstations boot from each server to prevent the control system from becoming disabled in the event of a server crash.

4 Conclusion

Network reliability can be defined as the ability of the network to remain operational (available) when some of its components have failed. Failure states in the network are caused by device failures, such as transceiver chatter and/or link failures. In general, these equipment failures are assumed to be statistically independent. At the APS we have provided a great deal of protection against single failures. In our experience, hardware failure has caused 95% of network crashes. To provide the reliability demanded by the APS users, we have provided a network with sufficient redundancy to meet the requirement of 99.95% availability.

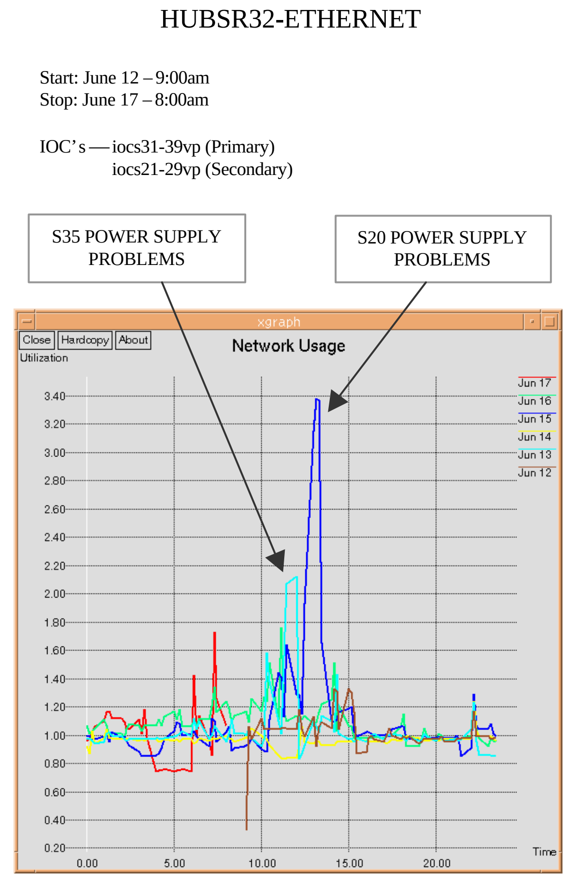
HUBSR32-ETHERNET

Start: June 12 – 9:00am
Stop: June 17 – 8:00am

IOC's — iocs31-39vp (Primary)
iocs21-29vp (Secondary)

S35 POWER SUPPLY PROBLEMS

S20 POWER SUPPLY PROBLEMS

Figure 4. Sniffer plot.