IEEE Transactions on Nuclear Science, Vol. NS-32, No. 5, October 1985

CONTROL SYSTEM FEATURES OF THE ARGONNE 6 GeV SYNCHROTRON LIGHT SOURCE*

M. Knott, G. Gunderson, F. Lenkszus, and W. McDowell
Argonne National Laboratory
9700 S. Cass Avenue
Argonne, IL 60439

Introduction

The Argonne 6 GeV synchrotron light source design consists of an electron/positron linac, a fast-cycling 6 GeV synchrotron, and the storage ring itself. The design attributes are presented elsewhere in this conference.[1] Three aspects of the overall design call for special attention in the control system design: First, the operation of a high energy positron accelerator in a fast cycling mode may demand high processing performance and high data throughput rates. Second, the high energy and small beam size projected (100 x 200 microns) will call for high resolution data processing and control precision in many areas. Finally, the necessity to provide independent, orthogonal control for each of up to 32 insertion device light beams both from the point of view of the experimental requirements and from the need to remove the effects of component vibration will require dedicated, high performance processors.

Control System Features

As with most modern accelerator control systems distributed processors will be used to perform in parallel the many specialized tasks that would be nearly impossible to accomplish in a single centralized facility. The processors will be specialized as to the nature of their task and chosen from the wide variety of equipment available commercially. At the highest level, one or more "host" processors implemented with superminicomputers will provide file storage, printing, archiving, alarm and situation displays, program development facilities, parameter logging, and high-speed, powerful computation service for the lower levels. At the next level, the operator consoles will be implemented with minicomputers specialized for graphic display generation and operator interaction. These computers will run most of the accelerator control and monitoring software, augmented by the host processor where needed.

At the next level closer to the accelerator is a minicomputer or microprocessor cluster, designated as a system computer, which is tailored to the task of translating the high-level needs of the console and host computers into the command streams to be directed to distributed single board computers (SBC's). Conversely, the system computer also translates the returning data into high-level responses and formats them for efficient transmission and use by the console and host computers. At the lowest, most distributed functional level are the SBC's. These processors are separated as to function (vacuum, magnet, diagnostic, etc.) within a cluster but share hardware interfaces and network access. A special type of distributed intelligence will be provided for each insertion device region to provide the high performance necessary for the high-speed, orthogonal compensation needed to maintain multiple light beam aiming points and angles.

*This work supported by the U. S. Department of Energy

The final, and possibly most important, element of any distributed computer system is the interconnecting network. If all of the processors in the system could directly access all data throughout the complex, an interconnecting network would be unnecessary. But lacking omnipotent processors and given the realities of time and distance, the choice and use of this network becomes an important issue.

Special Features

Although nothing in above description indicates radical departures from modern accelerator control system design, there are at least three areas where recent advances in computer equipment evolution have produced methods and devices that are particularly applicable to accelerator control system design. These developments are the computer-aided engineering (CAE) workstation, the very high performance SBC, and the local area network (LAN).

CAE Workstation-Based Consoles

The last several years have seen the rapid development of computer aided design and computer aided manufacturing (CAD/CAM) workstations. These systems demand a high performance computer closely coupled with a high resolution display and some form of graphic input or pointing device. The earliest form of CAD/CAM system employed a superminicomputer dedicated to a single user because of the computations required for the application and display tasks and the fast response desired. Now that microprocessors are attaining the required performance, either singly or working in concert, such systems are finding their way into other applications.

Accelerator control system consoles have developed similar features to CAD/CAM systems, namely dedicated console computers, graphic displays, pointing devices and the ability to control several separate display screens. These systems required large outlays in hardware and software development since most elements were not commercially available. We think that a point has been reached when most of the usual accelerator control console features can be provided by a generic implementation of the CAD/CAM system, the computer aided engineering or CAE system. Further, since most hardware and software tools are provided, the undertaking would be cost-effective.

Figure 1 is a physical diagram of a console position whose central feature is a CAE workstation. The touch-screens and alpha-numeric "comfort" displays are controlled by dedicated tasks running in the CAE computer. Communication with the host and system computers is by means of a high throughput LAN. Application tasks of low to medium complexity such as display updating, control knob servicing, and "virtual parameter" controls also run in this computer. Tasks that require sophisticated accelerator modeling or are in general highly compute-bound are relegated to the host computer with requests and results passed over the LAN.

0018-9499/85/1000-2014$01.00© 1985 IEEE

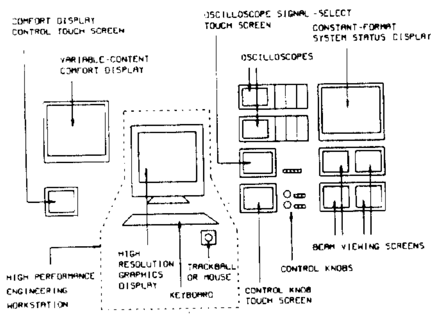

Fig.1: Major features of the central consoles

Figure 2 shows a typical display with several application programs sharing the high resolution display. Each program uses one or more display "windows" to interact with the operator visually. Pull-down menus are used by the computer and operator to communicate choices. The windows can be sized and moved by the operator as needed by the specific task being performed. A separate microprocessor provides the window management, and pan and zoom features. In this way, these time-consuming and sometimes complex tasks are off-loaded from the console computer. More importantly from the point of view of initial development cost and ease of future feature development, the software to accomplish these feats is off-loaded from the in-house staff.

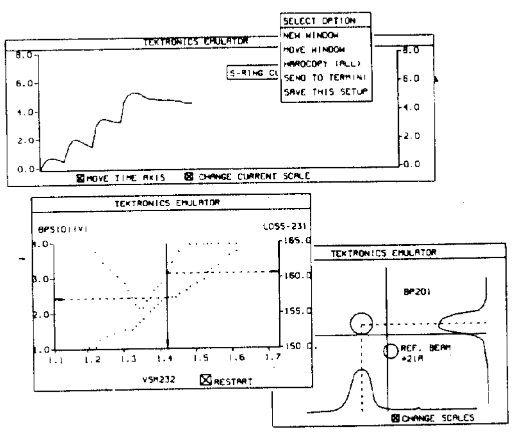

Fig.2: Typical multi-window display

High Performance SBC's

Single board computers will be used at the lower, machine interface level and these will be assigned responsibilities as functions, signal quantities and geography dictate. For the insertion device region, there is a need for a very high performance sensor/computation/control system. It is felt that a viable algorithm can be implemented to correct the positron beam position for element vibration and experimental needs without significant interference with downstream experiments. Since it is desired to correct for photon beam position errors as small as ten microns with a frequency content of up to 1 KHz, it will be necessary to measure the photon beam position in two planes at two locations, compute the necessary corrections and transmit these to the correction elements at a rate of about 10 KHz. It is clear that a high throughput, dedicated processor system will be needed.

Although it would seem that we have a natural application for an analog feedback loop, we feel that an all digital computation loop has the advantages of stable performance and arbitrary complexity (within ever-expanding limits). Although such an analog feedback loop has been designed and operated with one insertion device[2] and expanded to five such loops, these are based on three magnets and a single photon position detector in the horizontal plane. The system planned at Argonne will use four magnets and two photon beam position sensors in both planes. The four magnet control system gives the ability to control both angle and position but is more more complex than the essentially linear three magnet, single aiming point system. There are bit-slice processors and array processors on the market to solve the algorithm and separate SBC's could be assigned to tasks of measurement, beam location, and correction application, all tightly coupled and highly parallel in function. If interaction with adjacent loops were necessary, communication could be provided via parallel links. A LAN would be used to monitor and control the various loops. For example, all loop gains could be gradually adjusted to control turn-on and turn-off transients.

Local Area Networks

The choice of network topology, protocol and data rate can have a critical bearing on the final performance and operating features possible in the total system. We have limited our choices to peer-protocol bus networks because they offer the highest potential throughput for a given media, direct access to all nodes by any node, and the highest reliability since there is no dependence on active message relaying nodes. Of these three features, the second offers tremendous operational flexibility in that it allows any node (or a maintenance technician at that node) to access data at any other node on the network. "Broadcast" messages can communicate data or timing information to all nodes simultaneously. Contiguous parts of the network can operate in a stand-alone mode, a feature useful during construction and maintenance. In general, bus network systems allow the simulation of the omnipotent feature alluded to above.

Ethernet: On the surface, the network protocol with the best combination of features is that which employs carrier-sense, multiple-access with collision detection (CSMA/CD) such as Ethernet does. It has a bit rate of 10 MHz and virtually instantaneous minimum access to the media. However, since simultaneous use of the media by more than one node is impossible, any "collisions" which occur must be resolved and the protocol uses a random delay technique to do so. Figure 3 shows how the network remains stable with throughput growing linearly with an increasing number of users until about 90% of theoretical capacity.[3] It levels off at this point because of the increasing need to resolve collisions. This need to resolve collisions makes the maximum time to guarantee access to the media unpredictable, an undesirable feature in control systems.

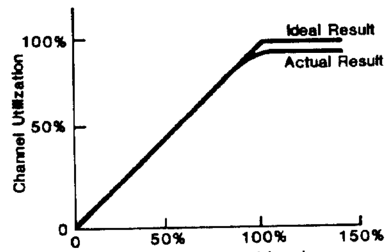

Fig.3: CSMA/CD throughput with increasing load

Token Passing Bus: A network protocol which avoids this problem is the token passing bus. In this method, permission to use the media (the token) is passed from node to node in a round-robin fashion. If the maximum message length is kept short enough, the maximum time to gain access to the media can be reasonable, but in any case the protocol is deterministic. As the number of nodes on the network grows however, this time also grows. Figure 4 shows how these two protocols compare as the number of nodes increases.

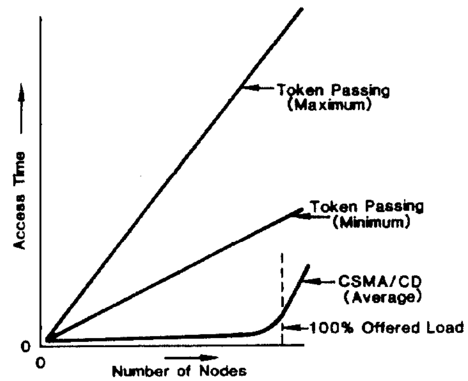

Fig.4: Access time comparison (arbitrary units)

Perhaps the critical factor in comparing these two protocols is the way in which the network will be used. An access method based on collision resolution tends to perform poorly if the nodes tend to communicate simultaneously. But, this is precisely what typically happens in an accelerator control system. Figure 5 shows how the access time for the CSMA/CD protocol increases more rapidly with increased number of users if the access attempts are synchronized than if they occur randomly. This is due to the time needed to resolve all of the "deliberate" collisions.

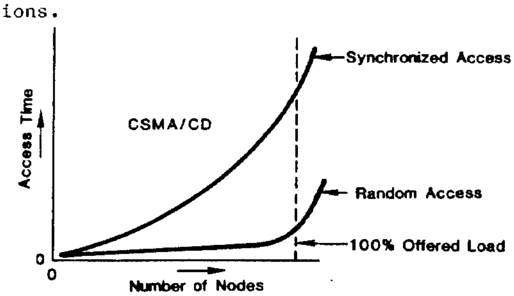

Fig.5: CSMA/CD access time growth

Another factor in assessing the performance of a network protocol is whether a targeted node can successfully handle the incoming traffic. A high performance network can overwhelm a node to which all or most of the traffic is addressed, again just the situation which occurs as synchronized readbacks funnel toward a system computer. The token bus protocol allows an overwhelmed node to hold the token until it frees enough buffers to handle an expected burst of messages. Another solution is represented by the "modified token passing" scheme used by the ARCNET protocol employed in some control system applications at FERMILAB.[4] In this system the receiving node is first queried as to the availability of an input buffer. While taking additional time, the loss of data and the complexity to recover the data is avoided. For these reasons, we are currently considering the token bus protocol the best choice for the several system networks.

The CSMA/CD or Ethernet protocol does perform well where the messages tend to be random in occurrence and comparatively long and so we think it will be a better method for the host/console/system computer network at the highest level. It is well supported by the vendors being considered for the host and CAE workstation computers. Vendor support is always an important factor when planning a large hardware/software development effort. Figure 6 shows how all of the elements of this proposed control system are interconnected.

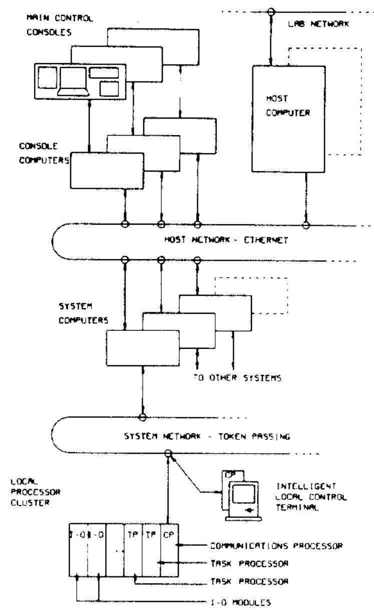

Fig.6: Major Elements of the Control System

Summary

As accelerators become more complex, their control systems are presented with an ever-increasing amount of data, precision, and parameter relationship complexity. Fortunately, the commercial proliferation of CAE workstations, the ever-increasing capability of microprocessors, and the advances made in multicomputer communication will enable accelerator control systems to keep pace with these needs.

References

1. E. A. Crosbie, et al, "Conceptual Design of the Argonne 6 GeV Synchrotron Light Source", to be published in the proceedings of this conference.

2. R. O. Hettel, "Beam Steering at the Stanford Synchrotron Radiation Laboratory", IEEE Trans. Nucl. Sci., Vol. NS-30, No. 4 (1983)

3. J. F. Shoch and J.A. Hupp, "Performance of an Ethernet Local Network -- A Preliminary Report," 20th IEEE Computer Society International Conf. (COMPCON-80, Spring) San Francisco, CA, Feb., (1980)

5. R. W. Goodwin and M. F. Shea, "Modern Control Techniques for Accelerators", Proceedings of the 10th International Conference on Cyclotrons and Their Applications, East Lansing, Mich., (1984)