

STATUS REPORT ON THE ADVANCED LIGHT SOURCE CONTROL SYSTEM

S. Magyary, M. Chin, M. Fahmie, H. Lancaster, P. Molinari, A. Robb, C. Timossi and J. Young

Accelerator and Fusion Research Division
Lawrence Berkeley Laboratory
University of California
Berkeley, California 94720
U.S.A.

Abstract

This paper is a status report on the ADVANCED LIGHT SOURCE (ALS) control system. The current status, performance data, and future plans will be discussed. Manpower, scheduling, and costs issues are addressed.

I. INTRODUCTION

The ALS control system was designed around the concepts of parallel processing, high CPU and I/O bandwidth, and human-friendly interface. Figure 1 shows the system architecture and its five primary layers (for details of the system see References [1] and [2]). Layer 1, represented by the Intelligent Local Controllers (ILCs), interfaces to the accelerator hardware and communicates with Layer 2, the Collector Micro Module (CMM). Layer 3 is the Display Micro Modules (DMM) that has bus access to the CMM and in turn communicates with the operator stations (Layer 4) via serial links. The operator stations are high-performance Personal Computers that have Ethernet network (Layer 5) access to file servers and other network services.

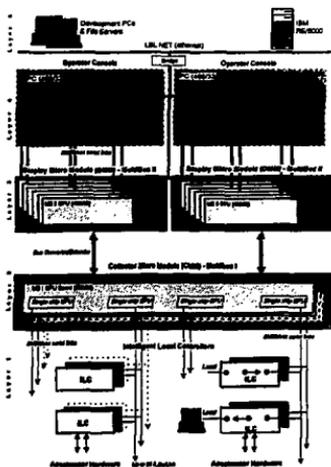


Figure 1. ALS control system architecture.

The ALS consists of an Electron Gun, a Linac, a Linac-to-Booster line, a Booster, a Booster-to-Storage-Ring line, a Storage Ring (SR), and a number of user beamlines. The control system is currently operating the existing parts of the ALS accelerator hardware consisting of the Gun, Linac, Linac-to-Booster line, and the Booster; the Booster-to-Storage-Ring line is being implemented now. The Storage Ring accelerator hardware is under construction; completion is expected sometime during the second quarter of 1992. We will then be ready to begin commissioning the SR via the control system, both locally and from the Control Room.

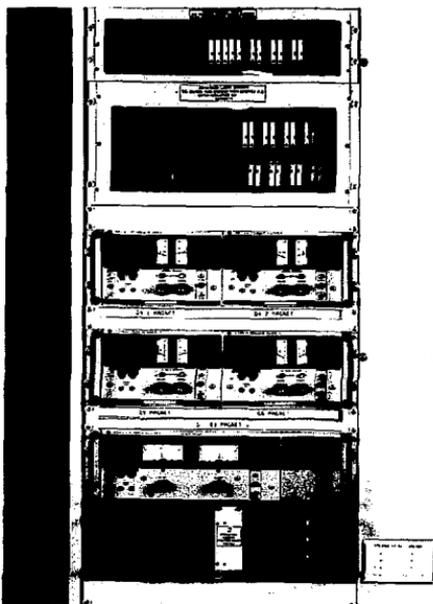


Figure 2. Typical ILC installation.

II. LAYER 1 (INTELLIGENT LOCAL CONTROLLERS)

The ILC is an intelligent controller consisting of an 80C186 main processor, an 80C187 math co-processor, and an 80C152 serial-control processor sharing 64 Kbytes of battery backed

memory. In addition, it has on board I/O resources of four 16-bit DACs, four 13-bit ADCs, 24 bits of digital control, and an SBX bus for expansion. It is a low power (< 5 watts), 3U high Eurocard-based controller in a shielded metal can that can communicate at a 2 Mbit/sec rate using twisted pair cabling. We had commercial companies build 200 of these first generation ILCs. Twenty of them (10%) were not functioning when received from the manufacturer; ten had minor problems (chip leads bent, missing chips, infant mortality, etc.) and were repairable. The remaining ten we have not attempted to repair (they have missing traces or shorts on the circuit boards) since we decided the cost of repair was not justified. We are currently using about 140 ILCs in the accelerator, with an additional 20 to be used for the Booster-to-Storage-Ring line. See Figure 2 for a typical installation showing 3 ILCs, Opto-22 interface, and the 3 quadrupole power supplies. We expect to use an additional 500 ILCs to complete the project (Storage Ring, Beamline Front Ends), however, these will be the next generation design. These new ILCs will have 16 Mhz, 80C186EB chip as the main processor (a 60% speed improvement over the older ILCs' CPU, and also lower-power dissipation), 256 Kbytes of memory, 16-bit ADCs and DACs, 16-bit SBX interface, a serial channel, and 28 Boolean lines. The analog and digital design is completed; layout and prototyping/testing will start shortly. The cost of the current ILCs is \$650 each; the new ones will be about \$950. To exercise the ILCs, and test the accelerator hardware, we use laptop computers (80386/80486-based), connected directly to the ILCs, using much of the same software as we use on the operating stations in the Control Room.

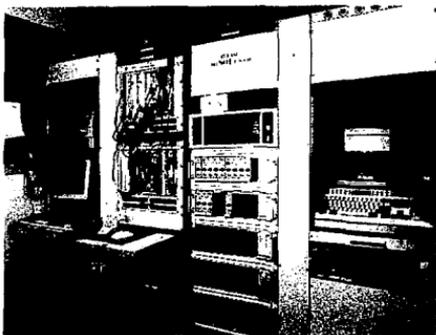


Figure 3. DMM, CMM, fiber optic interface, and file server.

III. LAYER 2. (COLLECTOR MICRO MODULE)

The CMM (Figure 3) contains all the data gathered from all the ILCs (i.e., it represents the entire accelerator database at any moment). The ILCs are connected to the CMM via fiber-optic lines. The serial communication on these lines is bi-directional (though using only a single fiber), and the bit rate is 2 Mbit/sec. We are currently using 12 of these lines for a total I/O physical bandwidth of 24 Mbit/sec. We use a commercial (Intel) 20 MHz, 386-based Multibus I board with 4 single chip processors (via custom built SBX modules) to service 4 serial lines. Therefore, to service the 12 lines currently in use, 3 Multibus I processors

are required. These boards allow us to handle approximately 800 messages, of about 75 bytes each, per second per line. Therefore, the total I/O bandwidth currently coming into the CMM is approximately 720 Kbytes/sec ($12 \times 800 \times 75$). This represents, on the average, about a 10 refresh/sec of the entire active part of the current accelerator database. We are currently evaluating a 33 MHz, 486 Multibus I board and are designing a new SBX interface that would allow us to service 8 lines per Multibus board at about 1600 messages/line/sec. At project completion, we plan to have approximately 64 lines (using 8 Multibus I processors operating in parallel) for a total I/O bandwidth of 128 Mbits/sec, and a useful data rate 7.5 MBytes/sec. This would represent, on the average, about a 15 to 20 per second refresh rate of the active database of the entire accelerator. We are also exploring the possibility of a second CMM; this would double our I/O bandwidth to 240 Mbit/sec at a modest cost.

IV. LAYER 3. (DISPLAY MICRO MODULE)

The DMM (Figure 3) consists of a 20-slot Multibus II System, bought as a unit from Intel Corp., that has fast parallel access (via a bus converter) to the CMM. This system contains a SCSI disc controller, an Ethernet interface, and a number of high performance singleboard computers. Under current operation, we are using 2 DMMs (we had not planned to install the second DMM till Storage Ring commissioning, so we are ahead of schedule in this area) to access the database in the CMM. The DMMs currently use RMX II as the real-time operating system (we are evaluating using RMX III, a full 32-bit system) using standard Intel hardware and software. The first DMM currently supports 6 (we have tried 7) commercial (Intel) 25 MHz, 486-cpu boards, while the second DMM currently uses 3 CPUs. The CPUs within each DMM operate in parallel, and each has 2 serial links (we promised 1, so this doubles I/O performance in this area) that directly connects it to the operator station. Each of these links has the same configuration as the links between the ILCs and the CMM (i.e., 2 Mbit/sec). At project completion, we will support at least 6 operator stations per DMM, for a total of 12 or more for the whole accelerator. We have enough bus bandwidth in the CMM so we could support even a third DMM. We plan to upgrade the CPUs in the DMM with processors that will be at least 2.5 times as fast as the current ones.

V. LAYER 4. (OPERATOR STATIONS)

The operator station is the human interface of the Control System; it presents accelerator data, including real-time data, scope traces, and live video (via multi-media programs and hardware) to the operator. The operator can use mice, keyboards, or nine dynamically assignable/labeled knobs. Six operator stations make up a console (Figure 4), and each console is supported by one DMM. The operator stations use 33 MHz, 486-based AST Personal Computers (PC), but they are upgradeable to faster CPUs via a simple card swap. Most use Windows 3.0 as the operating system, a few use OS/2 1.2. The serial communication links to the DMM allow us about 800-1000 database accesses/sec for each PC in the database client/server "request" mode. In the "driven" (i.e., when the DMM drives the PC) we can achieve about 1500 messages/sec. Both

of these rates are CPU limited, we expect them to be about 2500 and 3000 messages/sec (at those rates we will be I/O limited) respectively at project completion. Since we plan to have at least 12 operator stations, we will attain accesses in excess of 30,000/sec; however, even at these rates, we are using only about 10% of the bus and data bandwidth available in the CMM-DMM combination. These facts demonstrate the very high performance available in our star-based, shared memory approach to control system design. The choice of the PC as our human interface has proved to be a judicious one. As the popularity of Windows grows, the availability of commercial software for use in the control system grows rapidly. We currently make extensive use of DESIGNER (graphic editing tool), EXCEL (spread-sheet), TOOLBOOK (HyperCard-like package), VISUAL BASIC, TURBO PASCAL for WINDOWS, as well as the usual languages C, C++, etc. With these PCs, we also have available the usual collection of word processors, databases, and utilities (screen capture, etc.). The use of this commercial software has greatly increased our productivity, and has allowed us to keep our staffing requirements very low. The one area where we need improvement is in support of full 32-bit modeling applications. These are currently done on workstations and can access the database via Remote Procedure Calls (RPC). We plan to shift this work to OS/2-2.0, or its equivalent, as soon as possible. In the meantime, X-Window-based applications can appear as a window on the PCs under Windows 3.0 or OS/2, using commercial software.

A. Scheduling and Costs.

The control system is being used to commission the accelerator as it is being built. We are on target both in terms of cost and schedules, and are beginning to shift over to pre-operation. The staffing to date has been exactly as projected (with five programmers, one coordinator, and one half of an electronic designer) over the length of the construction project. Cost to date is approximately \$3.65M of a total projected cost of \$5M. Storage Ring commissioning is scheduled for the 2nd quarter of 1992, with project completion in 1993.

VII. ACKNOWLEDGMENT.

This work was supported by the Director, Office of Energy Research, Office of Basic Energy Sciences, Materials Sciences Division of the U.S. Department of Energy, under Contract No. DE-AC03-76SF00098

VIII. REFERENCES.

- [1] Magyary et al., "Advanced Light Source Control System," Lawrence Berkeley Laboratory, LBL 26028 and Proc. 1989 IEEE Particle Accelerator Conf., Chicago (IEEE Publishing, New York, 1989).
- [2] Magyary et al., "The Advanced Light Source Control System," Nuclear Instruments and Methods in Physics Research, A293 (1990) 36-43, North Holland.

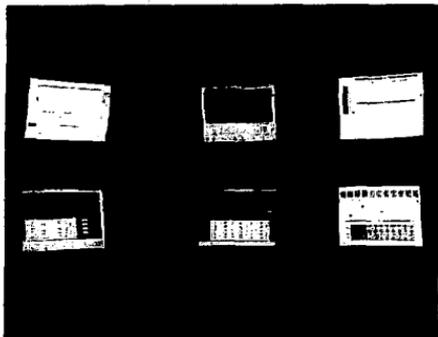


Figure 4. Operator console with six operator stations.

VI. LAYER 5. (NETWORK)

The PCs are networked via Ethernet to a Laboratory-wide network. This allows workstations, etc., to access the database via RPCs. An IBM RS6000 Workstation is used by the physicists for modeling and 32-bit numeric intensive applications. We also have a PC on this network with a 600-Mbyte disk as a file server for the operator stations and software development. A spare server is available to minimize down-time in case of failure by the main server. This network will also be the "user" interface into the control system for wiggler/undulator and beamline front end controls. A protection scheme to limit user control, to specified devices only, is under development.