

A HIGH PERFORMANCE CONTROL SYSTEM FOR A HEAVY ION MEDICAL ACCELERATOR

H. D. Lancaster, S. B. Magyary, and R. C. Sah
Lawrence Berkeley Laboratory
University of California
Berkeley, CA 94720

Summary

A high performance control system is being designed as part of a heavy ion medical accelerator. The accelerator will be a synchrotron dedicated to clinical and other biomedical uses of heavy ions, and it will deliver fully stripped ions at energies up to 800 MeV/nucleon. A key element in the design of an accelerator which will operate in a hospital environment is to provide a high performance control system. This control system will provide accelerator modeling to facilitate changes in operating mode, provide automatic beam tuning to simplify accelerator operations, and provide diagnostics to enhance reliability. The control system being designed utilizes many microcomputers operating in parallel to collect and transmit data; complex numerical computations are performed by a powerful minicomputer.

In order to provide the maximum operational flexibility, the Medical Accelerator control system will be capable of dealing with pulse-to-pulse changes in beam energy and ion species.

Control System Design Philosophy

The availability of vastly improved computer hardware at moderate cost means that it is now reasonable to consider building a computer system to fit an application. A very high performance computer system can be designed to provide monitoring and control of devices and instruments, as well as high-level control functions such as automatic beam-line tuning.

In particular, the reduced financial pressure to share computing resources among many functions has led to the concept of providing distributed intelligence. There is a clear trend to use larger numbers of computers in accelerator control systems: PEP used a few computers, the SuperHILAC Third Injector used about twenty, and ALS will use over two hundred. One important advantage of using distributed intelligence is the dramatic simplification of the software.

To simplify the operation of the Medical Accelerator, a large number of interactive, easily-understood graphics displays with simple control functions will be provided. These displays and the supporting calculations will require extremely large computing power to be adequately responsive, on the order of 0.1 second. To achieve these responses, we will use several high performance microcomputers performing tasks in parallel and communicating via a multi-processing-bus. This concept, which replaces minicomputers, was developed and demonstrated on the SuperHILAC Third Injector control system, and will be used throughout the control system when high computing rates are needed. (1,2,3)

Reliability and maintainability will be stressed in the control system design. This applies to its

This work was supported by the U.S. Department of Energy under Contract No. DE-AC03-76SF00098, and in part by the National Cancer Institute of the National Institutes of Health under grant No. CA-19138.

own operation and its ability to aid in the repair of other accelerator components. With this in mind we will stress modularity in component design for fast system repair by replacement, and adequate monitoring of accelerator component parameters so proper operation can be determined by the control system.

As much as possible, risks associated with technical development will be minimized by using familiar design concepts and technology. We will remain open to new advances in technology by designing into the Medical Accelerator control system the flexibility to incorporate new technology as it becomes available.

High-level control functions, such as accelerator modeling, consist primarily of number crunching rather than data manipulation. What is required then, is a medium-sized computer. The critical element here is a friendly environment for software development, because software costs will dominate hardware costs. Communications requirements between this computer and the rest of the control system are not particularly severe, because only moderate quantities of data need to be transferred.

Control System Architecture (Figure 1)

This control system is characterized by the use of many dedicated computers each of which performs a definite, fixed function. The two fundamental units of this system are the "microcomputer module" and the "intelligent local controller" (ILC).

Microcomputer Module

The microcomputer module consists of a card cage with a Multibus* backplane. This card cage contains a number of boards, at least two of which are single board computers which incorporate high performance microprocessors. A key element in this design is the use of local busses on the computer boards, because then the Multibus can be reserved solely for communications between boards. True parallel processing is achieved. Incidentally, the selection of the Multibus standard is a distinct advantage because of the many vendors who produce Multibus-compatible boards.

There are three types of microcomputer modules, each serving a different function:

1. The Display Microcomputer Module (DMM) services the operator console. Access to the CMM data base is achieved by using "Multibus extension boards" which permit the DMM to address directly the CMM data base. A word of data can be accessed in about 2 microseconds. There is one DMM for each operator station, but more operator stations can be provided by adding DMM's. The DMM also is the gateway to the Ethernet link.

The DMM consists of a card cage with three Intel iAPX 286 computer boards, and they perform a number

*Multibus is a trademark of Intel Corporation

of functions. The "operating computer" executes standard operating tasks, refreshes CRT display pages and provides a gateway to Ethernet. The "console computer" keeps track of device settings which are entered using the knobs and touch panels at the operator station. The "graphics computer" provides support for graphic displays. Other boards in the DMM include color graphics display boards, the Multibus extension board which provides communications to the CMM, and an Ethernet control board. The gateway to the Ethernet link will provide for access to a network resource manager (File system, Disc, spooling printer, etc.) for temporary and permanent housekeeping functions (logging, etc.). These functions are sophisticated in performance and slow in average link data rates. They are therefore decoupled from the fast response control system by the Ethernet network.

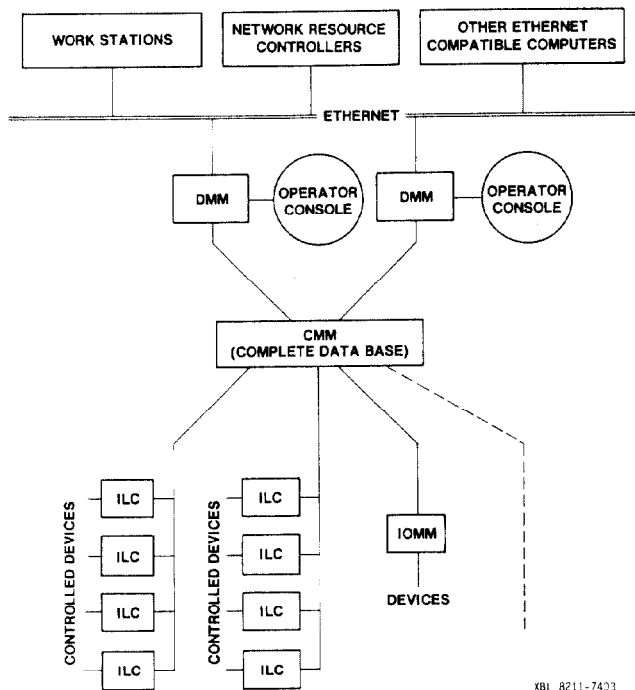


Figure 1: Control System Architecture

2. The Collector Microcomputer Module (CMM) is the single module which collects the data from all of the IOMM's and ILC's. Data are transmitted over links at a 180 kilobaud rate. The CMM has a data base resident on dual-ported RAM, and this data base includes copies of all the local data bases.

A very rough estimate of the magnitude of the monitoring and control task for the Medical Accelerator indicates that possibly 400 devices, with 3000 signals, would have to be controlled. This corresponds to about 20 IOMM's or 400 ILC's or an appropriate mixture. There appears to be no difficulty with the size of the data base, when the CMM accommodates 20 IOMM's. If we assume that each IOMM has a 12 kilobyte data base (both active and passive portions), then 20 IOMM's would require a data base of about 0.24 Megabytes. Since most new microprocessors have an address space of 16 Megabyte, the DMM address space is large enough to include the CMM data base plus its own Multibus requirement.

The communications boards will be intelligent, containing single-chip microcomputers. Each single-chip microcomputer talks directly to the communications channel from an IOMM (or from the ILC's which might replace the IOMM), and each

microcomputer deals with the portion of the data base corresponding to its IOMM. If we were to put 4 microcomputers (along with dual-ported RAM for the data base) on each board, then 5 boards could handle 20 IOMM's. An example of a single-chip microcomputer we can use the Intel 8751, which has an integral serial communications channel. This design avoids the sharing of a resource (the Multibus in this case) and eliminates a potential bottleneck.

The CMM contains a Multibus extension board, which permits the DMM to access the CMM data base and several communications boards for linking to IOMM's and ILC's.

3. The Input Output Microcomputer Module (IOMM) provides an interface to the accelerator equipment by collecting data and transmitting control instructions. The IOMM is used where parallel computation is required to achieve the desired throughput. The IOMM contains a local data base. There are at least two computer boards in the IOMM, one of which controls input/output to the devices, and one of which deals with the data base and messages. Other boards handle the actual input/output of analog and boolean signals. One IOMM services about 20 devices, such as magnet power supplies, with a total of about 150 signals (20 ADC, 20 DAC, about 110 boolean). The functions performed by the IOMM include data collection, device control, closed loop control, data base, and timing. The IOMM uses the information in its local data base to control devices, and new data are collected to refresh the entries in both the local data base and also in the complete data base at the CMM. The software for the IOMM's consists only of applications programs which are stored on EEPROM's or EPROM's.

Intelligent Local Controller

The Intelligent Local Controller (ILC) is a single board computer with serial communication, a high performance microprocessor, and sufficient input-output to control and monitor an accelerator component. The ILC contains a local data base.

The Medical Accelerator must be a facility which provides routine treatments within a hospital environment. As such, it must meet stringent requirements of reliability and maintainability. Furthermore, the need to minimize operating costs by keeping the operating staff very small also implies very high reliability and maintainability.

First, we note that overall reliability and maintainability would be greatly enhanced, if the control system could monitor sufficient signals. For example, it would clearly speed repairs and minimize down-time, if the operator were instantly informed of the reasons why power supplies tripped off.

A second important point is that having the simplest possible software at the IOMM level maximizes the ease of trouble shooting. Although the IOMM applications programs are already very simple, some unavoidable complexity remains because of the necessity to monitor and control about 20 devices.

By following our basic control system design philosophy, we can see that an attractive solution is to distribute when possible the IOMM tasks among a large number of dedicated computers. That is, we propose to replace each IOMM with ILC's, each of which is located at the single controlled device for which the ILC is responsible. By essentially eliminating the cabling problems inherent in transmitting signals between controlled devices and

IOMM's this strategy makes it convenient to monitor a larger number of signals at each controlled device. Furthermore, the greatly simplified software at each ILC would significantly enhance the ease of trouble shooting malfunctions at this level of the control system. The increased modularity of the hardware would also tend to make system expansion easier.

Although it may become possible to purchase an appropriate ILC board within a few years, right now we conceive of an ILC as a single, specially-built board. The ILC would include a microprocessor (perhaps an Intel 80188), DAC, ADC, PROM, RAM, and serial transmitter/receiver. The functions performed by an ILC would include all of the IOMM functions plus added monitoring and improved diagnostics (perhaps a self-test mode). EPROM-resident programs are anticipated to be the most convenient option for ILC's, since ILC software, being even more modular, is likely to undergo even fewer changes than IOMM software.

Communications

An important consideration in the design of a control system is to organize the interactions between the different levels of computing. An extremely simple solution to this problem will be implemented. The IOMM and CMM function to produce an up-to-date data base at the CMM. The CMM serves as memory for the DMM.

The communications between the CMM and the DMM depend upon two "Multibus extension boards", which allow the computer at the DMM to gain access to the data base at the CMM by simply addressing the appropriate portion of the data base memory.

There is a single, dedicated communications channel between each IOMM and the CMM. The CMM database is updated continuously. A much smaller volume of control data goes from the CMM to IOMM.

If an IOMM is replaced with ILC's, then the communications problem with the CMM becomes more complex. That is, each of the ILC's must now communicate with the CMM, as compared with one IOMM previously. As before, the traffic on the communications link would consist primarily of fresh data being sent to the CMM, with a considerably smaller volume of device control data being sent to the ILC's.

To handle this communications problem a dedicated channel will be used to link the ILC's with a microprocessor at the CMM. A "master" would be required to control traffic on this channel. A very simple scheme of traffic control, such as using dedicated time slots, will be used and a general "network" would not be required. The reason is simply that ILC's would not be given the capability to communicate directly with one another. That is, if two devices (such as two magnet power supplies) must interact with one another, they do not "talk" directly to each other. Instead, fresh data concerning these devices are collected at the data base at the CMM, and the applications program at the DMM, which has access to the data base, provides the desired interaction.

Number Cruncher

An additional computer, linked to the DMM's, via Ethernet is required in order to provide a substantial number crunching capability in the Medical Accelerator control system. The primary purpose of this computer is to achieve a certain

degree of automation. By performing high-level control functions, this computer can reduce the number of actions and decisions which are required of the accelerator operator. At a minimum, the computer can give instructions or enumerate options, but the goal is to approach fully automatic operation. Automated control functions will include accelerator modeling (to provide guidance for changes in operating parameters), automatic tuning of ion sources and beam lines, diagnosis of accelerator malfunctions, and closed orbit corrections. For the most part, we anticipate that these programs would need to be run on an infrequent basis, mainly during accelerator tune-up.

The data link between the number cruncher and the rest of the control system presents no special problem, because the required data rates are expected to be very modest. Certainly an Ethernet communications link can provide adequate speed.

The more general programming and computing requirements for the number cruncher differ considerably from those of other computers in the control system, so it appears attractive to think of this computer as a separate entity, set somewhat off to the side.

References

1. H. D. Lancaster et al., "A Microcomputer Control System for the SuperHILAC Third Injector", Proceedings of the 1979 Linear Accel. Conf., BNL 51134 (1979), 269-273 (Also LBL-9763).
2. S. Magvary et al., "A High Performance/Low Cost Accelerator Control System", IEEE Trans. on Nucl. Sci. NS-28, No. 2 (1981), 1461-1464 (Also LBL-11761).
3. S. Magvary et al., "Operating Experience With A New Accelerator Control System Based upon Microprocessors", IEEE Trans. on Nucl. Sci. NS-28, No. 3 (1981), 2201-2203 (Also LBL-11724).