# TRENDS IN ACCELERATOR CONTROL SYSTEMS

Rolland P. Johnson

Maxwell Laboratories, Inc.
Brobeck Division, 4905 Central Avenue, Richmond, CA 94804

*Abstract*

A survey of control systems at national labs shows that in-house, bottoms-up developments have had varying degrees of success, largely dependent on the level of effort invested. New laboratories and projects continue to develop their own systems, although some collaborations have been initiated. Smaller laboratories with limited resources are not likely to find satisfying solutions by emulating the larger labs. A new trend is underway whereby software and hardware experience and developments from larger national labs are incorporated in carefully planned systems for smaller facilities.

## 1. INTRODUCTION

With a few exceptions, most accelerator control systems have evolved to a similar architecture. Namely, to a system of three levels of computers; data acquisition, central server, and operator's console, all connected by some local area network.

The only significant deviation from this architecture is with shared memory approaches (ALS, SSC, and CEBAF) where technological advances are prerequisites for the transmission of large amounts of data to a central computer. This is not unlike schemes of 15 or 20 years ago when microprocessors were less convenient. The problems of bandwidth encountered in these approaches are being attacked in the hope of simplifying the access to data by higher-level processes and thereby reducing the software development effort. It remains to be seen whether the technology is up to the task and whether other costs are reduced.

Indeed, except for these exceptional approaches the major concern of the moment is to minimize the costs of software development. This shift of emphasis over the last 15 years is a consequence of the present availability of powerful, inexpensive computers and the increasingly high cost of sophisticated manpower needed for specialized applications. It is not unusual to see estimates of control system costs which show software to be more than 80% of the total.

One can imagine two trends developing. One is the use of commercial software being adapted to the accelerator environment. The other is the sharing of software among different projects and laboratories.

### *1.1 Commercial Software*

Adoption of commercial software to the accelerator environment is an attractive prospect. The uses of database products like ORACLE and SYBASE are certain to reduce costs, though it often takes a clever design to incorporate these products without causing a serious degradation of real-time performance. Other products, such as spreadsheets like EXCEL and 2020, are also being used to take the place of accelerator control applications that have traditionally been produced by local experts.

I have used both approaches to solve the problem of ramping control of a synchrotron (a fortran scheme for Tevatron control and 2020 for the LSU synchrotron) and find advantages to each. The most extreme example of the use of commercial software for an accelerator system that I have seen is probably the Isolde Control System. It will be interesting to see how such a PC-based system will do for a more complex machine with difficult real time requirements.

### *1.2 Software Sharing*

Software standards and world-wide sharing of software, while not yet a trend, is a question of some importance. The present situation is that the developments in different labs are difficult to transpose. At the console level, programs have been exchanged. The lack of documentation, as is common in national labs, has made this a difficult business, however. At lower levels, the sharing of software seems almost impossible because there are rarely adequate standards even within one laboratory.

## 2. REBUILD, FROM THE BOTTOM UP

For the most part, early control systems developed from the bottom up. That is, the concerns were first for the individual devices in the tunnel, only later for groups of devices, and even much later, if at all, for the accelerator or accelerators as a whole. The first problem was to control power supplies which were distributed over a large area. Thus the first control systems were not much more than remote control systems for power supplies to change currents and turn them on and off. Additional functions to alarm on trips or incorrect readbacks and to allow saving or restoring values were next to come. Once diagnostic devices like beam current and position monitors were digitized and available to the control system a new level of sophistication was possible. Namely programs could be written which could significantly aid the accelerator physicist to make the machines run well. More recent developments, especially considering the increased power of the control computers, include rather complete models of machines which can be included in the on-line code to make these higher level programs even more effective.

### *2.1 New Wheels, Old Problems*

One of the trends in particle accelerator control systems that I do not want to talk about is that which is a reiteration of the learning curve which was just described. There are still places where the emphasis is to build a control system, bottom up, out of high tech components with very little thought as to what the system is supposed to do and how it will evolve with machine development and further technological advances. There is the continuing trend of mistaking a control system for a data acquisition system. There are still people using buzz-acronyms to talk about control systems, without understanding the implications of such choices on the expense and complexity of the system.

These traditional trends are accompanied by traditional pains when the new control system is commissioned with the new accelerator. This is usually exacerbated by having insufficient manpower to do the program development and documentation tools, forcing machine physicists to become professional programmers or vice versa. Invariably there is too much to do and the machine commissioning has to be done with an incomplete control system. And as anyone knows who has ever tried to make a machine work for the first time with any number of important people looking on with skeptical anticipation, that is when you need all the help you can get.

### *2.2 Software Expense*

Having seen the development of control requirements at several accelerators and at different laboratories, one could anticipate it and come up with a better plan for a new laboratory or project. For example, one would normally pay most attention to those aspects of a control system which are the most expensive. Just to make life interesting, I suppose, there are several ways that a control program can be expensive. To the head of a controls group, the number of man-months of programmer effort may be the big expense. To the head of the operations group it may be expensive to not have a program that could have diagnosed why his accelerator ran with unacceptable intensities for several days. To a laboratory director with an annual budget of a few tenths of a billion dollars, a few days of machine studies to make a program work properly may completely overshadow the man-months of programmer time.

As you can guess, I believe these higher-level applications programs are very important. They are the last things to be developed. Often they incorporate the hard-won lessons of the machines in question and they are necessary to make the machine work well. They also tend to exist for the life of the machine, unlike the computers, for example, which are usually replaced after 6 or 8 years as unmaintainable.

At the very largest and now oldest accelerator laboratories such as Fermilab, CERN, and SLAC, there was never any question but that the machine control systems should be developed along with the machines they controlled. Indeed, often each machine at each of those laboratories was developed with its own rather independent set of controls modules, diagnostics, and computer programs. The concepts were new, as were the machines, and in most cases the control systems grew like topsy, one machine having little influence on another, even at the same laboratory.

### *2.3 Unified Controls*

As these laboratories matured it was generally recognized that there was much to be gained by having a unified system for machine control at a particular laboratory, and several plans were undertaken to achieve this goal. To the best of my knowledge, the first and perhaps only example of a successful implementation of a completely unified control system for a major accelerator complex was that initiated with the construction of the Tevatron over ten years ago. All the accelerators and beam lines at Fermilab, old and new, are now controlled from the same small control room using the same consoles using programs which, if not identical, use the same conventions for operator interfaces. Program code is developed and maintained in common and there is a single machine access database. Physicists, engineers, and programmers are able to work on any machine more or less interchangeably.

And at Fermilab we have started to see programs developed which are used on more than one machine, even for those which are traditionally very machine-specific. The mechanism for this phenomenon is rather undefined. Sometimes a programmer is asked to reproduce a program he has written for a second machine and finds it possible and easier to generalize the program. An engineer responsible for the BPM or flying wire systems decides that he wants the programs to be general and useful on any machine. Basically, whenever a good idea comes along or a good program developed, there is a demand that it be made available to other machines. Since each machine uses the same control system, it is particularly easy to develop a clone. But the future trends will involve developing the programs as general-purpose solutions to a common problem right from the start. In any case, there is now a growing body of programs which could be quickly adapted to any synchrotron.

## 3. COST EFFECTIVE ANALYSIS

In the United States there are three large labs which seem to be following the old traditions of starting from scratch with a bottoms up approach (APS, SSC, and CEBAF). One assumes these enterprises will succeed because sufficient resources will be applied. Other laboratories have a problem to solve which can be attacked in a new way which I would like to suggest is going to be a future trend.

The main feature of a future approach to a control system is a cost-effectiveness analysis of the entire system. All components listed below will be included in the equation.

### *3.1 System Services*

Reliability, accuracy, response speed, fast graphics.
Alarms and Limits, Save and Restore, Data Logger.
Fast time plots, On-line Documentation, hardcopies.
Engineering level support, e.g. parameter pages.
Integrated clock and timing system.
Well-tested database system.

### *3.2 Program Development Tools*

System Maintained, code captured.
Large library of routines.
Debugging facilities, documentation.

### *3.3 High Level Applications*

Often unique to a particular machine.
Often require valuable machine time to develop.
Usually cost several times the hardware.
Should survive changes in any hardware.
Ready at earliest days of machine commissioning.

### *3.4 Hardware*

Important, but viewed as temporary.
Select most cost-effective solution integrated 5-7 years (mean time between replacements for computers).
MTBR slightly longer for control modules.
Controls requirements different than data acquisition.
Standardization, support, more important than MIPs.
Modern architecture uses distributed intelligence for speed, flexibility, economy.
Choose devices which minimize software development effort.

### *3.5 Hardware Maintenance*

Software system for hardware, network maintenance.
Microprocessor module support.

## 4. NEW TRENDS

Having listed all the components in a control system with some comments on the criteria for selecting them, we come to the real question which may lead to the next trend. Namely, how might a smaller or less experienced laboratory acquire all of these things in a reasonable time and with an acceptable expense? The only possible answer, of course, is to adopt a system from another laboratory. I will ignore the possibility of considering commercial systems used for process control as a solution to the problem as they invariably do not supply enough of the required features.

There have been examples of laboratories adopting the controls schemes from other places which have had various degrees of success. The touch-panel driven Nodal system developed at the CERN SPS during the seventies was adopted at DESY. More recently, the CEBAF system has been tried by the XSLS at Brookhaven and the refrigeration system at the SSCL in Dallas.

At the present time, the Brobeck Division of Maxwell Labs has entered into an agreement with

Fermilab to develop the accelerator control system used for all of the Fermilab machines into a commercial product. The goal is to produce a generic control system for any sized accelerator laboratory from something as large as the SSC to as small as the most modest synchrotron light source or proton therapy synchrotron. The plan is to start with the programs and accelerator-specific hardware developed at Fermilab and further develop them for use at other laboratories. Maxwell Labs will then provide the customer with a complete control system which has been tailored to his specific needs and contains all the functionality discussed above.

This is a new concept for accelerators, of course. The idea that you can get an integrated control system that will do what you need from the start. That such a thing is possible is a reflection of the experience of the operation of dozens of machines which, in the end, have had the same problems to solve. That it is desirable to just buy what you need follows from the very large expense needed to create a control system.

One expects that a laboratory with an "off the shelf" control system will be spared the lost time, money, and manpower associated with system development. For smaller laboratories with limited resources, to buy such a system may be only way to acquire something satisfactory. It is also an excellent way to provide an infusion of expertise, providing training and guidance to a neophyte staff.

I thank many colleagues at CERN, DESY, FERMILAB, and MLI for stimulating conversations.