Who Gets to Specify the Control System?*

by Elliott McCrory†*, Jean-Francois Ostiguy and Shekhar Mishra*
Fermi National Accelerator Laboratory
Batavia, Illinois, 60510, USA

1. Introduction

The users of an accelerator control system are typically considered to be: 1. the operations experts, 2. the machine engineers, 3. the accelerator scientists and 4. the administration of the facility. Each of these groups places different demands on a control system, each (except possibly the administrators) putting an equivalent load on the overall system. The operator needs a reliable and simple system every minute, but the demand at any time is not very high. The engineer and the scientist may need a lot of information, but only at very specific times, depending on the nature of the situation. This paper considers the demands on the control system from each of these groups, and how the needs of one group may tend to override the needs of another.

Firstly, we discuss the historical context in which control systems have developed, pointing out where things have changed significantly. Secondly, the four groups of workers which can potentially place requirements on the control system design are defined and their demands are considered. The interplay of the sometimes-conflicting demands of these groups is discussed. Thirdly, using Fermilab's recently completed 22-month run as an example, the amount of time spent in various aspects of a run is derived. Finally, some conclusions are drawn and recommendations are made.

A preliminary definition of terms is necessary. A "run" is an extended operating period of the complex in which no major changes are made to the complex itself. At Fermilab, a run may last for many months. The "complex" refers to the collection of all the components in the overall facility. A "component" includes everything which makes up the complex, from entire accelerators (the Tevatron) down to the small pieces of apparatus (a trim magnet). The "Controls Group" is the administrative organization whose primary function is to design, build, commission, maintain and improve the hardware and software in the complex which allows a worker (or an automated computer system) to view the operation of the complex from a location which is possibly removed from the place at which the observables are measured, for example, the Control Room.

2. Historical Perspective

The accelerator control systems of the 1970's evolved out of the work of the individual machine groups in the Fermilab complex. The Controls Group was formed in the 1980's in an attempt to integrate all of these various and independent control systems together into an "Accelerator Control System." The primary goal in this effort was to unify the control system in view of the disparate machines. This was an in-house effort, with no external commercial software, hardware or expertise.

Throughout this period, operating the complex was the overriding requirement for the control systems. It was necessary to derive paradigms for operations so that the operator, engineer and scientist did not have to learn a different paradigm for each accelerator or for each component. This effort has been extremely successful.

In the 1990's, the situation has changed substantially. In particular, the Accelerator Controls Groups are no longer on the leading edge of developing computer control technologies. This distinction has moved to Industry so that the newest technologies are no longer being developed for science.

---

* *Work supported by the US Department of Energy under contract number DE-AC02-76H03000*

† *Presenter*

It is a fact that over the last twenty years, computer technologies (CPUs, networks, software tools, etc.) have doubled in performance every three years. There is every indication that this trend will continue. Nevertheless, accelerator control systems remain as good examples of excellently-integrated and stable computer systems.

Our facilities, and consequently our control systems, are under great pressure to improve even though our resources are shrinking. In this light, it is reasonable to take a step back and rethink what demands are being put on our complexes and, possibly, restructure our efforts to match.

3. Who are Potentially the Specifiers?

Who are those who can potentially make specifications for the control system? In this paper, we define those who can make these specifications as: the operations experts, the machine engineers, the accelerator scientists and the administrators. These groups represent a partially orthogonal set of concerns which spans the space of possible requirements. An individual worker can be in more than one of these classifications, and, in principle, can be in all four. The following sections define these types of workers in the environment of an accelerator complex. The definitions are based on how people are organized at Fermilab.

3.1 Operations Experts

These workers are responsible for keeping the complex running 24 hours per day, for a large part of year. They are typically on rotating shifts and are usually the youngest and most enthusiastic workers at the facility. They are trained to do routine monitoring and tuning of the complex and its components. They also can repair a large class of problems which arise, but typically, the solutions to these problems are well-defined, like how to change a faulty piece of equipment. Often, it is not possible for the operations expert to fix a problem, so he/she must know when to call an expert.

The operations expert needs to have a broad view of the entire complex, but depth of knowledge of any individual component is not necessary. Also, his/her temporal view is limited: How can this problem be solved today (before my shift is over)? Can we do something so this problem does no recur?

The operations expert demands a stable and consistent view of the complex through the control system. He/she needs application programs which inform him/her of problems ("Comfort Displays"), which automate tedious, repetitive and/or intricate operations ("Sequencer"), and which make all data available at any time ("Parameter Page").

A partial list of the types of data which the operations expert requires includes: Beam currents, particle intensities, ground currents, luminosities, beam emittance and component status.

3.2 Machine Engineers

Machine engineers are the technicians, engineers and scientists who build, maintain and fix specific components in the complex. They usually are not on shift, but may be on call 24-hours per day. In an accelerator, these components include sub-accelerators (like the Main Ring or Booster), the magnet systems (the magnets, the power supplies and the connections between them), the RF systems, the cryogenics systems, the beam diagnostics and the vacuum systems, to name a few. These workers are required to fix their equipment when it is broken, write embedded-system software, write some component-specific application programs, implement improvements in the components and assist in the design of new components.

The machine expert needs to have a clear, uninterrupted view of his/her component. His/her attention span is shorter and longer than that of operations: fixing a piece of equipment takes an hour, but maintaining it takes months or years. The time span for improvements is longer still.

The machine expert demands to have total access to the internals of his/her component, often when standing in front of that component, but also from his/her office or the Control Room. He/she also needs high flexibility during the implementation phase of his/her component. Rapid prototyping of their equipment through a graphical display is desired.

The machine expert needs the same type of data as the operations expert, but his/her view tends to be more concerned with trends in these parameters.

3.3 Accelerator Scientists

Accelerator scientists are the engineers and scientists who integrate new features into the complex and who are required to diagnose and subsequently fix the really hard problems which occur during the course of the run. They are also required to understand and verify the theoretical limits of the complex, to specify future direction for the complex, to suggest software and hardware improvements and to design new components. To do these things, accelerator scientists are often required to perform experiments on the complex or on specific components in the complex.

An accelerator scientist needs to have a clear view of the entire complex, as does the operations expert, and, from time-to-time, must also have the myopic perspective of the machine expert. He/she has an attention span of the operations expert ("fix it today and for the run") and, additionally, somewhat longer, too ("how can we fundamentally improve *this* in the long term?")

The accelerator scientist demands time-correlated data. In fact, he/she demands that the control system behaves like a data acquisition system. But, primarily, he/she must have an extreme level of flexibility in controls, adding new, possibly temporary hardware and software to the system in order to diagnose a specific problem. He/she needs to be able to get data for a different type of computer in order to solve unusual problems (e.g., putting a PC into a UNIX-oriented console environment).

The type of information needed by the accelerator scientist includes: emittance growth, luminosity lifetime, integrated luminosity and general performance limitation.

3.4 Administrators of the Complex

This group of workers is included for completeness, since the demands which they place on the control system are small. In this definition, the administrators of the complex listen to the other three groups of workers and, based on the data presented from these groups (data which is usually obtained through the control system), they decide the future direction of the facility. Their need for a continuous, reliable overview of the status of the complex is obtained, often through great effort, through a semi-static, site-wide display. At Fermilab, this is realized through a closed-circuit TV channel 13.

4. Interplay of Demands

To summarize, the operations expert demands a stable, consistent view of the complex so that problems in the complex can be easily identified. The machine expert demands to have a clear view of his/her piece of equipment and needs a lot of flexibility to choose the latest and most modern tools within that equipment. The accelerator scientist demands flexibility, especially during commissioning, and good data acquisition functionality.

The key area in which these demands collide is the operations experts demand for consistency and the machine experts and the accelerator scientists demand for flexibility. In the past, this conflict has been "resolved" by saying that any of the demands of the machine experts and the accelerator scientists must be implemented in a way that does not violate the operations experts demand for consistency. This has all-too-often meant that the specific needs of the machine expert and the accelerator scientist have not been met because a canonical solution cannot be found. In particular, there has been a reticence to allow many commercial software packages to be included in the control system because they are typically difficult to integrate fully into the existing system. This makes rapid prototyping and "temporary" software for experiments rather difficult to implement. For example, if an machine expert is having difficulty finishing the installation of his piece of equipment, he/she needs to be able to use whatever tools are at his/her disposal to solve the problems. If these tools include software which is not fully integrated into the control system, then it may be difficult to use the software, and, at the worst, he/she cannot get the needed information at all.

5. How is Time Spent During a Run?

What fraction of the time is *really* spent in operations? In order to get an accurate breakdown of how time might be spent during a run, we call upon our own direct and recent experience with the just-completed run of the Fermilab Collider "Run 1B." It began in September of 1993 with the commissioning of the 400 MeV Linac and ended in June, 1995:a total of 22 months. The goal of this run was to increase the integrated luminosity delivered to the experiments each month by about a factor of two over the best performance of the previous run. Moreover, we were scheduled to operate the complex for about twice as long as in the previous run, so the experiments were expecting a 4-X increase in the number of events to tape. This run had a long period of very successful operations, and, in fact, we delivered record luminosities for an extended period, delivering a total of 5 times more integrated luminosity than in the previous run. However, most of the time was not spent "running."

The main feature of the time line, shown in Figure 1, is the extended periods in which we were had severe problems. The most significant of these was "interval 3" in the Figure. During this time, the Tevatron was incapable of delivering reasonable luminosities because the transverse coupling in the accelerator was quite high. This led to enormous emittance growth and severe current limitations. After several months of intense detective work, including new diagnostic software and hardware, it was determined that a low-beta quadrupole was displaced by a large fraction of a millimeter. The magnet was centered, but it was rolled. Correcting this roll fixed all the problems in a dramatic fashion! The other period, marked "interval 6" in the Figure, was caused by a piece of wire in the Main Ring, which limited the intensity and emittance in the transfers through that accelerator.

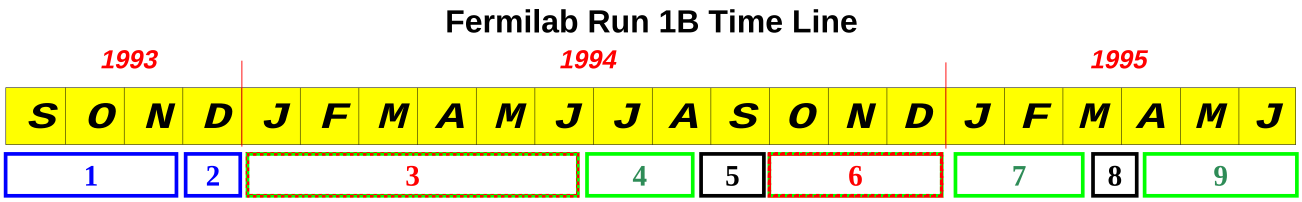

1. Commissioning of the 400 MeV Linac, the refurbished Booster, Main Ring and PBar Source
2. Commissioning the Tevatron
3. Attempt routine running, but encountered severe problems (in Tevatron) which reduced the effectiveness of the period significantly (see text).
4. Located problem in Tevatron, fixed it and had a period of good, uninterrupted running
5. LN2 Supplier decides to send LN2 to McDonalds instead of Fermilab. Had planned a shutdown for Sept/Oct, so did this shutdown anyway, unprepared.
6. Recovery from LN2 loss, severe problems in Main Ring makes this period less than optimum
7. Problems solved, good running period.
8. Planned M&D Shutdown to install new Main Ring Coalescing cavities.
9. Good running period.

Total time Run 1B: 22 Months = 4 + 3 + 6 + 9.

Figure 1., Fermilab Run 1B time line.

In summary, the time spent during this run was:

Commissioning & Improvements	4 months (18%)
Shutdown & Recovery	3 months (14%)
Run with Severe Problems	6 months (27%)
Operations	9 months (41%)

The distinction between *Operations* and *Commissioning & Improvements* (which includes experimental studies on the complex) is difficult to determine. Here, we use the observation that during the best weeks (of which there were 15), we were able to obtain more than 3 inverse picobarns integrated luminosity at each of the experiments. This normalizes to 13 pb^{-1} per calendar month. Since we provided an integrated 115 pb^{-1} for the run, that would mean that we had the equivalent of about 8.8 months of "good operations" for the run. The rest of the time was spent either explicitly doing non-operational things (shutdown, recovery, beam experiments) or getting up to the 3 pb^{-1} level. Our best two weeks had an integrated luminosity greater than 4.5 pb^{-1}, so even weeks of 3 pb^{-1} had failures and beam studies.*

In these monetarily uncertain times in the USA, the demands placed on our complex by outside forces (e.g., Congress) are increasing while the amount of money we get is decreasing. For example, Fermilab is trying to have the $130 million Main Injector accelerator ready for "Run 2" in 1998. It is likely that this accelerator will be commissioned by the smallest staff the Accelerator Division has had in twenty years, if present trends continue (509 FTE staff in 1982; 458 FTE now; no hiring). Moreover, the time between then and now will be packed with lots of non-operational activities. Thus, the relative lack of operations will continue, and probably get more pronounced, into the next millennium.

6. Conclusions from this Temporal Analysis

The majority of time at a complex like the Fermilab Collider is *not* spent in routine operations. New components are being commissioned for each run (by definition). Thus, we need to become more proficient at commissioning components at all levels. Machine experts must be able to use the most modern equipment, without excessive concern about how to incorporate their component into the Control System. Accelerator scientists must be able to invent new and innovative software methods easily and quickly, and be able to obtain large, time-correlated structures of data from components so that the new components can be commissioned efficiently. It is time to fundamentally incorporate these requirements into the Control System, possibly at the expense of the "stable and consistent view of the complex" which the operations experts have demanded.

7. Implications

In short, the following specifications are the important ones for today's control systems:

1. Control Systems must be flexible enough to adapt to the quickly-changing complex,

2. There needs to be as small an overhead as possible for adding new hardware and new software to the control system,

3. The data acquisition capabilities of the control system need to be enhanced, including:

    3a. More data acquisition bandwidth

    3b. Fundamentally available time stamp on all data

---

* One can recast this in terms of what we *thought* we were doing at the time, and we would have:

Commissioning, shutdown, recovery & general problems	10 months (45%)
Run with Severe Problems	2 months (9%)
Operations	10 months (45%)

4. Commercial products, both hardware and software, should be easy to incorporate into the control system,

5. Computer hardware should be easy to change in order to take advantage of the latest technology; this would apply to the console workstations, front end crates and networks.

6. Applications software should be capable of running under different operating systems. This would be possible if there existed a well-defined application program interface (API) for a control system. Decoupling this API from network and hardware details is essential.

7. Operations experts must accept that their view of the complex is not the only one−their view must coexist with other views within the control system.

8. Specific Recommendations

We are not in a position to make lasting, hard-and-fast specific recommendations on what changes must be made since today’s recommendations are out of date tomorrow. Some structural recommendations can be made. It is critical for the Controls Groups to forcefully introduce new computer technologies into service at the earliest possible time. A mature controls group in the present budgetary context tends to not recruit new blood. This makes it even more important to take positive action to ensure that the personnel remain up-to-date and motivated.

There are many promising new features in Control Systems out there which deserve mention. In software and protocols, cdev may work out−the general idea of making a standard controls API is good. TCL/TK and it’s derived packages are a very promising, robust systems for doing the windowing interface for all facilities. In fact, TCL/TK have been ported to PCs, so this becomes a platform-independent way to do the GUI. There is a trend towards POSIX-compliant operating systems, which is an excellent basis for guiding future software portability. Our documentation problems may have been solved by the excellent HTML tools ("Web Browsers") available today. And, object-oriented techniques have enormous potential for mitigating controls software complexity at all levels. Recall a nice paper from 91-ICALEPCS [ref: Cork & Nishimure, "Framework for Control System Development"], although many of its ideas are now dated.

In the realm of equipment-level processing, consolidation on the VxWorks operating system has been a great aid to the effort here. We need to continue to use POSIX-compliant OS’s, especially the POSIX.4 real time OS definition. There are two nice, forward-looking front-end ideas at Fermilab which deserve attention: MOOC (Minimal Object-oriented Communication) helps the front-end programmer deal with complexity through OO techniques; and the Open Access Front End allows scientist and engineer to measure, model or simulate any aspect of the complex in a robust and convenient fashion by making virtual "front ends" which can calculate a wide range of scientific and engineering quantities on line.

9. Conclusions

Today, we are faced with a complex which is constantly changing. The rate of this change is increasing, but our human resources are shrinking. Also, the users’ idea of "capable control system" is becoming more sophisticated. Moreover, the computer-related industries have orders of magnitude more resources for producing stable products than any of us can claim. Therefore, it is necessary to understand and cope with the following consequences: The primary goal of the design of a control system must be flexibility; Control systems must include good data-acquisition capabilities; Commercial products usually should be chosen over home-brewed applications; The view of the complex by operations is going to be less stable.