What constitutes a good control system for operations?

M. Lamont
CERN, CH-1211 Geneva 23, Switzerland.

Abstract

The requirements that operations have of a control system are enumerated without reference to implementation. Examples from the SPS and LEP are used to illustrate possible solutions and pitfalls. The key constituents of a control system from an operational viewpoint are then highlighted. Some suggestions about how they might be provided are discussed in the conclusion.

1 Introduction

One seemingly eternal debate in recent years is on high level software to be used by operators and machine physicists to run an accelerator. Here we step in the murky area between the control system and the people who use it. In simple terms, most accelerators are built to deliver some sort of beam to physicists. Operations' job is to run the machine from day to day, from year to year, to deliver this beam. Operations interacts with the accelerator via the control system. The control system, as used by the operators, should therefore be designed to enable them do their job as effectively as possible [1]. This is not always the case.

In what follows a general overview of operations' requirements is sketched. Some key areas of importance are highlighted. Drawing on experience with LEP and the SPS, examples are given where these requirements have been met, together with some in which things are far from satisfactory. Based on this, the key constituents of a good control system are enumerated. Some indications of how these might be acquired are suggested.

The main concern here is the high level control system. It is assumed that the middle and lower layers allow the fast and reliable exchange of data and provide the requisite functionality, such as synchronization via a timing system and so on. Clearly, however, inadequacies at these levels are reflected in the performance of the higher level.

2 What is Operations' job?

Operations are required to run an accelerator, usually on a 24 hourly basis, to deliver high quality beams in a controlled and predictable way. They are also intimately involved in machine development, in exploring the potentiality of the machine and in exploiting news ways of running it.

There is a large amount of different equipment to be controlled, complicated procedures to be followed and a large amount of data from beam instrumentation and the equipment which is used to monitor the condition of the beam and the machine. Each of the many subsystems of an accelerator has its own vagaries and pitfalls which have to be dealt with.

In addition there is usually pressure to optimize the performance of the machine and break last year's record peals luminosity, integrated luminosity, protons on target and so on.

Operations also perform a key coordinating role in ensuring that all groups are aware of the ongoing requirements of the machine.

3 Requirements

Given the above job description and based on the experience of the SPS and LEP operations group, there follows an attempt to enumerate the general requirements of operations:

1. To be able to drive the machine through it's duty cycle in an organized, efficient and reproducible way. In normal physics operations reproducibility is the key, however, the flexibility must exist to be able to drive the machine in different ways during machine development.

2. For the settings associated with a wide variety of equipment to be properly managed. This is of fundamental importance to the efficient operation of a machine. In a super-cycling machine like the PS or SPS the required

functionality includes the ability to construct so-called super-cycles from various cycles used to inject, ramp and extract leptons or hadrons. In LEP such things as the ability to switch between different optics and machine configuration are important. In general, facilities are required that allow retrieval, roll-back and modification. For all equipment the settings should be in a common repository.

3. To be able to control the state of the equipment. One should be able to address one piece of hardware, a group of hardware or the whole machine.

4. To be able to diagnose, and rectify as far as possible, equipment faults.

5. The ability to control the machine in terms of relevant parameters e.g. tune, chromaticity or the closed orbit. Clearly this is one of the most fundamental manipulations. Any changes should be fast and reversible even faster. There should be a history and the ability to step back to an arbitrary point in time. The interfaces should be generic; it should be possible to do everything in the same way as far as possible .

6. Ability to perform post-run analysis. All key machine parameters should be logged at appropriate frequencies. A mechanism is required for retrieval and display.

7. Measure and correct. Acquire a measurement on request, process it and present it to the user in a meaningful way. There is the optional need to accept input, perform calculations e.g. orbit correction, and possibly send a correction to hardware.

8. The possibility to develop complex tastes involving multiple trims and measurements. For example, a scan in which a parameter is varied and associated measurements made at each stage.

9. Fast beam diagnostics such as online display of lifetimes, spill structure, backgrounds.

10. Fixed displays. The state of the machine at a glance.

11. It should world reliably.

12. It should world fast.

13. It should world now.

The above requirements need to be met in an environment that provides:

- Good communication with the experiments and other accelerators.
- Standard tools for data display.
- A console manager.
- Standardization of look and feel.
- An integrated alarm system.
- Standard error handling.
- Standard HELP facility.
- Control system diagnostics.
- A minimum number of operating systems.

- A reboot facility.

4 LEP and SPS - some good points

The SPS has been running for over 20 years, LEP for about 6 years. Although there are still some shortcomings in the control systems of both machines, there's a lot to be thankful for. In an attempt to highlight some of the good constituents a closer look is taken at some features of both systems.

4.1 Trim parameter

The ability to set a current in a magnet is not considered sufficient.

Operations are only rarely interested in setting the current in an individual magnet. The principal concern is that the relevant power converter is supplying the requisite current and even then we are only interested when it is not. To reiterate we are not concerned with low level details of specific machine components except when things go wrong.

When we change parameters in the machine we do so in terms of:

- Physics parameters: tune, chromaticity, B-field.
- Combination of magnet strengths calculated off-line e.g. coupling compensation using the skew quadrupoles.
- Quantities involving online calculation: bumps, orbit corrections, separation.

There is a need for these trims to be recorded in the same place, use the same mechanism to apply them, have the same history mechanism and be in the same place for post run analysis. In both SPS and LEP analysis and design led to the recognition of the similarity in trimming any given parameter e.g. chromaticity:

1. delta Q'_h input by operator,
2. algorithmic calculation of required strength change in the sextupoles,
3. conversion of new strengths to current,
4. send the new current values to the equipment,

and, say, the synchrotron tune:

1. delta Q. input by operator,
2. algorithmic calculation of total RF voltage,
3. conversion of total voltage to individual RF unit voltages,
4. send the new voltage values to the equipment.

Implementation resulted in a single application or function call to trim any parameter in the machine. All parameters are treated in the same way and thus bookkeeping software is written once and then comes for free. This is good in terms of development cost and good for operations because everything is treated in a standard way.

For the SPS, the resulting application allows the trimming of the functions on a given cycle in terms of physics parameters, hardware magnitudes or hardware settings. In LEP there are two applications: one for the trimming of the ramp and squeeze, the other for steady-state trims which is used, for example, when the machine is in coast. This concentration of functionality did not come free, however, the effort is well rewarded, it works and it is appreciated.

4.2 Measurements

Rationalization of the acquisition and treatment of measurements has been a source of continual frustration. Why is it so difficult?

1. The wide range of data structures that beam instrumentation can come up with.
2. The sheer amount of data that can be generated.

3. Individual creativity.

Attempts have been made to find generic solutions to this problem. They have inevitably, in our experience, not been able to encompass everything. A well written dedicated application would sometimes seem to be the best solution for operations. This software could well be developed by the instrument group themselves but guidelines and tools must be in place to avoid chaos. Some things that have worked in this regime for SPS and LEP are outlined below.

4.2.1 THE DATAVIEWER

One solution to one part of the problem seems to be to provide a common data viewing tool to be used by many applications. In SPS and LEP the same tool, the so-called dataviewer, is used. Good examples of its use in LEP are the orbit correction package and the Q-meter application, which present widely varying data with very different operator input using the same dataviewer. The functionality of the dataviewer was provided after extensive dialog with operations. The advantages are: it is a standard tool and thus familiar to the operators, it is powerful with all the required functionality e.g. save, print, zoom, trim, input etc. and it involves lower application development cost.

4.2.2 MEASUREMENT DATABASE

In LEP and SPS an attempt has been made to use an online database as a common repository for measurement data. It has proved extremely useful but again is not universally applicable [4]. The volume of data can just be too high. However, it has allowed the development of tools such as a generic fixed display program and a postrun analysis tool. Together with trim parameter functionality it allows the development of quite complicated procedures. e.g. Vernier scans (adjustment of the collision point of beams using small changes to electrostatic separators): Tell the experiments, trim separators, measure luminosities, beam sizes and separator voltages, loop and then plot the results. This was possible in about 300 lines of code. One could, of course, argue that programming should not be necessary but it is a very powerful alternative. The advantages of having all the data, measurement and control accessible in the same way cannot be over stressed.

4.3 Other good points

In the SPS and LEP, high level control system settings generation and management is probably one area where things are more-or-less healthy. Both machines make use of an online database: C-tree for the SPS and ORACLE for LEP [2, 3]. The databases were well designed after considerable analysis of the requirements. The software was designed to provide appropriate functionality. It works well. All equipment functions for the SPS main ring and LEP are contained on their respective database. A database might not appear as such a high priority for steady state machines but the data robustness has proved invaluable for both the SPS and LEP.

The vagaries of the different equipment are encapsulated in what is known as a black box. All black boxes are callable in exactly the same way. This has allowed rapid development of standard interfaces to the equipment and allows everything to be driven in essentially the same way.

Some crucial beam related information, such as lifetimes and beam currents, short cut the control system proper and make use of video links from the low level crates. This has proved invaluable in LEP when fighting beam-beam driven lifetime problems. The information also makes a more stately progress through regular channels.

5 Some bad points

Both the size of the SPS and LEP and some peculiarities of their control system architecture means that they are not particularly fast. This is reflected mainly in slow transfer of large amounts of data, for example, in closed orbit acquisition.

A lot of the software has been developed by operations themselves and sometimes pressure of work leads to less than exacting effort in clearing-bugs.

There are at least 5 different groups producing software for the control room: controls, operations, beam instrumentation, beam transfer and the RF group. This leads almost inevitably to proliferation of applications, solutions and the flowering of individual creativity. Often a problem is solved more than once. Ad hoc solutions are provided by people busy with other concerns, the hardware for instance, and sometimes a decent interface is a while coming.

In the case of LEP, some seemingly important functionality was late. For example, it was 3 years before we had an online display of the beam-beam tune shifts, one of the key parameters used to tune the luminosity.

Successful coherent solutions have been overtaken by upgrades which has led to platform dependent solutions. The control system for the SPS and LEP now spans: OS9, Xenix, Lynx/OS, DEC, NODAL, HP/UX, Apollo, DOS and even Windows. There is also a variety of communication methods and buses, including RPCs, sockets, milview-1553, BITBUS and GPIB. Clearly providing a coherent solution in such an environment is more of a challenge.

6 What constitutes a good control system for operations?

Somewhat tautologically, one that meets the requirements of operations. In both SPS and LEP the most positive results have come from attempts to first establish the requirements of operations before providing a solution. Methods have proved invaluable and have been successfully used on more wide ranging projects. The following general points may be highlighted:

6.1 Applications

- Key functionality. To reiterate: run control (a sequencer, perhaps), settings management including generation, trim facilities, equipment diagnostics, logging, sensible treatment of measurements, ability to build complex procedures, fixed displays.

- Appropriate functionality. Only the required functionality should be presented to the operator. Unnecessary complexity should be shunned when designing interfaces for operations, one should at least tip one's hat in the direction of ergonomics. Expert actions should be buried but available.

- Coherent design which provides the operator with a small number of powerful applications.

6.2 Environment

- Standard tools

- Same look and feel to all applications, but no good getting too fixated.

- Common help and error reporting.

- A decent console manager.

- Online databases.

- The usual adjectives: reliable, fast.

- A common operating system as far as possible.

- A good reboot facility.

7 Conclusions

> . . . when we recognize the battle against chaos, mess, and unmastered complexity as one of computing science's major callings, we must admit that 'Beauty is our Business'. - Edsger W. Dijkstra

The constituents of a good control system are fairly easily enumerated. Operations requires data management, data visualization, equipment control and monitoring in a standard environment. However, the complexity and diversity of accelerator systems pose a severe challenge in providing such facilities.

It is not enough to provide equipment access and a toolkit. The software and data management requirements have to be properly analyzed. Coherent functionality can only come from such an analysis and subsequent design. This is, after all, a discipline. The same is true of Object-Oriented programming. One still has to understand the problem and adopt a method. Objects are not a panacea. Here I can but cite the Object-Orientation analysis gurus: Martin, O'Dell, Grady Booch, Shlaer and Mellor.

Pragmatism and appropriateness is the key. One should not get fixated on a method. Smaller applications can get by without it after asking operations what they need.

In a diverse environment, management is also important if a lasting solution is to be found. While individual groups and individuals are free to pursue their own vision, coherence will be hard to achieve. Reinvention and the not invented here syndrome clearly costs a lot in wasted effort. At the end of the day it is the operators and the efficiency of the machine that suffers. This is clearly more of a problem with large groups.

Finally, in all this, one should remember-that the machine performance is the priority.

References

[1] Control systems of accelerators, P.Wolstenholme, Proc. Cern Accelerator School 1984, CERN 85-19 (1985).

[2] R.Bailey, A.Bell¬, P. Collier, M . Lamont , G .de Rijk, M .Tarrant , Development of the LEP high level control software using ORACLE as an on-line database. Nucl. Inst. and Meth. A347(1994) 430-433.

[3] R.Bailey, P.Collier, M. Lamont, Performance of the high level application software during LEP operation. Nucl. Inst. and Meth. A347(1994) 350-353.

[4] M. Lamont, The measurement system. Proc. fourth workshop on LEP performance. CERN SL/94-06(DI) (1994) .