The Impact of New Accelerator Control Software on LEP Performance

R. Bailey, A. Belk, P. Collier, M. Lamont, G. de Rijk and M. Tarrant
CERN

CH-1211 Geneva 23, Switzerland

Abstract

After the first year of running LEP, it became apparent that a new generation of application software would be required for efficient long term exploitation of the accelerator. In response to this need, a suite of accelerator control software has been developed, which is new both in style and functionality. During 1992 this software has been extensively used for driving LEP in many different operational modes, which include several different optics, polarisation runs at different energies and 8 bunch operation with Pretzels. The software has performed well and has undoubtedly enhanced the efficiency of accelerator operations. In particular the turnaround time has been significantly reduced, giving an increase of around 20% in the integrated luminosity for the year. Furthermore the software has made the accelerator accessible to less experienced operators. After outlining the development strategy, the overall functionality and performance of the software is discussed, with particular emphasis on improvements in operating efficiency. Some evaluation of the performance and reliability of ORACLE as an on-line database is also given.

1. Brief history of the project

In response to requests from many areas, not least the LEP operators themselves, during the summer of 1990 the authors began considering what was required of application software for routine operation of LEP. Later in the year and into 1991, these ideas were developed, resulting in a rather detailed structured analysis of the required functionality by spring 1991.

The data flow diagrams and data dictionary of this analysis formed the basis of software design and data modelling work which was undertaken between March and September 1991. In parallel with the latter stages of this, database requirements were investigated which resulted in a definition of the structure of the on-line data and the choice of ORACLE as the on-line database for driving LEP.

We were thus in a position from autumn 1991, after one year of analysis and design work, to undertake module coding and database implementation. This activity of course continued through implementation and testing ready for first use for operations in spring 1992.

2. The major tasks performed

The software is divided broadly into four distinct parts, each interacting directly with the on-line database. These software packages are either driven through interfaces running under the standard SPS/LEP console manager [1], or through commands from other software. This last facility allows the operator to control the run by using a task sequencer to drive the new software. Before discussing these four pieces of software it is necessary to understand something of the underlying data organisation.

2.1 Data organisation

The accelerator is divided into around 30 different systems. Within each of these systems the data is organised in a hierarchy, consisting of physics parameters at the highest level, descending through hardware magnitudes (which in the case of magnetic systems is the field strength) to hardware settings such as currents. It is these hardware settings that are loaded into the machine. For the most part the operator wants to interact with the machine in terms of physics parameters, with propagation down through the hierarchy to the hardware settings being taken care of by the software. Data for each system is described as amplitude versus time functions, reflecting the need to take LEP from injection energy, through an energy ramp to physics energy and through a vertical beta* squeeze to the physics optics.

The multitude of hardware elements are also grouped into around 20 different element groups. This makes it easier to interact with the hardware independently of any physics interest, such as during start-up when one wants to load lots of equipment, or during periods of equipment problems. The way in which the operator can interact with any piece of equipment depends on the equipment type (power converter, kicker etc.), and this interaction is completely data driven from the database.

2.2 Settings Generation

The overall task performed here is to take the output of the MAD program and to build settings in a suitable format for running the machine. That is to say, this software produces all the data that the other applications require. The operator selects which energy ramp and beta* squeeze he wants to be used to make up a LEP run. Once having

CM-P00061223

CERN LIBRARIES, GENEVA

defined this, he can create on the database a specification that describes how he wants to run the machine. This run specification is used to create the full set of functions that are needed to effect the mode of operation required. A major part of this is to combine the ramp and squeeze functions to produce functions that define how the machine parameters vary from the injection point right through to the physics settings. These data form the full image of the machine.

In order to work more efficiently on the machine when in a static state, we have introduced the notion of a time slice of the machine settings. The most frequent examples of this are the settings for accumulation or for physics, but should the operator choose to introduce a breakpoint at which to stop in the middle of the ramp or squeeze for whatever reason, a corresponding set of machine settings is automatically produced.

2.3 Drive Hardware

This software provides a uniform way of sending all requisite equipment settings to the relevant hardware. There are two basic areas of applicability; either to load settings to the vast majority of equipment, or to load setting changes to the hardware group affected by a trim. This functionality is possible because of the structure of the data, allowing one to drive physics systems, hardware groups or individual pieces of equipment.

2.4 Trim Actual Settings

At a given breakpoint a time-slice of the active machine functions exists, called the actual settings. These can be trimmed either in terms of physics parameters or at the hardware magnitude or hardware setting level. Again a key feature is a unified interface for making trims to any system and at any level.

While making trims at a certain breakpoint, a complete history of the changes made is kept, enabling the operator to go back to some previous state of the machine. Before leaving the breakpoint the operator has the option of incorporating the changes that he has made into the functions. There is now considerable versatility in the way in which this incorporation is made. When the incorporation is finished, the actual setting trim history is archived and all actual setting trims removed.

2.5 Trim Functions

This offers the ability to trim physics parameters, hardware magnitudes of hardware settings defined as functions of time in the machine settings. This is presently done manually through a single graphical interface, but we have foreseen the possibility to feedback measured parameters, thereby providing an auto trim

facility. The same program, without the interface, is used to incorporate the actual trims made at a breakpoint.

In much the same way as for actual settings trimming, a complete history of all function trims is kept, enabling the operator to back out to some previous point in time. This is true whether the trim is made by hand, by auto trim or by incorporate.

3. Software performance during operation

The major applications described, and others, communicate with each other via the on-line database. They are all essentially data driven, enabling new systems or equipment groups to be added without writing new code. For these reasons the structure of the data is of utmost importance, not least in terms of the performance of the software.

3.1 Choice of ORACLE

There are many databases available, most of them commercially. Why choose ORACLE, which has little or no history of on-line implementation ? There are several reasons;

- Complete set of integrated software tools
- Available on a wide variety of platforms
- Professional support available
- Full RDBMS available
- High level query language, SQL

Apart from these considerations we were mostly worried about performance. Several benchmark tests were made, such as the time taken to perform a chromaticity trim, requiring some 20 functions to be modified on the database. These tests showed that for the frequently used applications, the response time was the order of seconds which is considered to be acceptable. Other pieces of software, such as generating the reference settings, can take up to a minute but since they are only run at the beginning of an operational period lasting weeks, this is not a problem. Indeed throughout the whole of the year, the database performed well enough for the needs of operations. It also proved to be extremely reliable.

3.2 Settings management

LEP was required to perform in a variety of different ways during 1992. For routine operation several different optics were used early in the year before the optimum one was established. Each new optics meant generating new settings (ramps and squeezes) and optimising them with beam. During this the new software performed well, leading to a significant reduction of the amount of time required to commission a new optics. Towards the end of the year this was achieved in a few hours whereas in 1991

it always took the best part of a day.

Furthermore, once a ramp had been commissioned it was much easier to keep track of any variation in the machine with time. Such variations led in previous years to a reduction in the efficiency of accelerating beam from the injection energy to physics energy, a figure of 75% being a typical weekly average. In 1992 this figure was often over 90%, see figure 1.

Figure 1 - Intensities at 20 and 45 GeV

The time spent going from physics conditions, through cycling the machine, accumulation of beam current, ramping and squeezing and back into physics also plays an important part in the LEP performance. In 1991 the mean time for this was over 3 hours for the 150 coasts made. In 1992 this was reduced to 2 hours 10 minutes over the 200 coasts made (Figure 2). In other words some 200 hours of physics time, out of 1700 achieved, was gained through this improvement.

Figure 2 - filling time in 1991 and 1992

3.3 Machine tuning and run analysis

The majority of running in 1992 was with a low emittance optics to maximise the luminosity. The price to pay for this was that the machine parameters had to very closely monitored and controlled throughout a physics run, which lasted typically around 10 hours. Particularly during the first hour or two of physics the operations team would be required to make many changes to key physics parameters such as beam sizes, betatron tunes and the closed orbit. All these changes could be achieved through a standard

interface, allowing quick and clear actions on the machine.

Many of the changes made were empirical trims, and the facility to revert back to any previous machine state, provided by the trim history, proved extremely useful.

In the longer term the trim histories proved valuable in post-analysis of different runs. It was possible to correlate changes in beam behaviour with trims made on the machine. It should be stressed that the choice of ORACLE again proved to be justified here, enabling these correlations to be developed interactively on the database.

3.4 MMI

Finally a word on the ease of use of the applications. In previous years, routine operation of LEP had been performed essentially by the engineer on shift. It had proved very difficult to delegate tasks to technicians, partly because of the complexity of operation but also because of the way the information was presented to him. With the new software in use during 1992, less experienced operators have felt much more comfortable and have been able to work effectively on the machine.

4. Conclusions

The new applications software for driving LEP has made an important contribution to the operation of the machine during 1992. The estimated gain in integrated luminosity is of order 20%. The look and feel of the software has made the accelerator more accessible to less experienced operators.

5. References

- [1] Experience with workstations for accelerator control at the CERN SPS
A. Ogle et al, Proceedings of Accelerator Controls Conference, Vancouver 1989