IEEE Transactions on Nuclear Science, Vol. NS-28, No. 3, June 1981

OPERATIONAL EXPERIENCE USING THE SPS CONTROL SYSTEM

G. Shering and V. Hatton
CERN, 1211 Geneva 23, Switzerland

The SPS control system introduced many new concepts and created substantial interest during its design and construction. It was the subject of papers at this conference in 1975 and 1977. In 1975 the design concepts were presented, then in 1977 the realisation and initial operating experience were described. Since then the SPS has had four years of continuous operation and development, up to the long shutdown in 1980 for preparation of its part-time use as a proton-antiproton collider. This paper describes, from the point of view of the Operations Group, the experience in using the control system, and its evolution, over this period.

It was quickly realised that the control programs determine the quality of the control system as seen in operation. Thus this paper is partly the story of the control or 'applications' programs, and their evolution. One of the features of the SPS control system is that applications programming is done by the users, and here the Operations Group have played an important part.

Other important parts of the control system directly affecting operation are the main control room and its facilities, the surveillance and alarm system, and the viewing of waveforms and analog signals. The performance and evolution of these aspects is also described.

SPS Operation History

In April 1976 10 GeV protons were injected from the CPS down the injection line TT10, and in June protons were accelerated to 400 GeV. Extraction followed soon after, and intensities of 5 * 1012 protons per pulse were achieved. By the end of 1976 the SPS was in regular operation for physics experiments in the West Hall at 200 GeV, and for neutrino physics at 400 GeV. At the end of the first four year operation period which is the subject of this paper, some 20 experiments were scheduled, sharing up to 2.7 * 1013 protons per pulse at a 12 second repetition rate. This represented about 85% of the beam from the CPS, and was supplied for about 90% of the scheduled time.

The SPS was about 3 years behind FNAL in its physics program, but had the advantage that an experimental area and several experiments were all ready to go in 1976. As a result there was strong pressure from the experimental physicists for good and reliable running right from the start. This in turn put pressure on the control system and our use of it. Fortunately, the control system had been considered an important and integral part of the SPS, and was immediately able to fulfill its role of complete computer control. Nevertheless this operational pressure on a new system led to a very rapid evolution in the first few years.

Operation Requirements

The control system is used in three main phases of SPS operation; start-up, regular running with optimisation, and machine development. There is a danger of considering a control system to be applicable to only one of these phases. In the SPS we initially put most emphasis on the commissioning of hardware and the start-up of the accelerator as these had to be successful to satisfy the physics requirements and to prove the new control concepts. Most of the publicity for the SPS control system so far has been linked to this first stage. This paper indicates how the SPS controls have evolved since 1976 in all three phases, and how the control system has made a vital contribution to overall accelerator management.

Although start-up is the shortest phase in normal operation, it is an exacting time for both the staff and the control system. We have certainly not reached the limit of computer assistance in this area of accelerator management.

Normal running lasts for about 10 days after start-up and is therefore the longest phase of operation. The most spectacular programs come into play at this time for machine optimisation. The surveillance and alarm programs are also vital during this phase.

The role of the programs during the machine development phase has varied over the years. Initially many programs were written specifically for machine development and have subsequently passed into operation. More recently this has been less frequent as the increasing sophistication of the operational programs has been more than adequate for most studies, except where new hardware has been introduced. With many new programs for anti-proton storage, however, the cycle seems to be starting again, although this time the operations staff are taking a much more important role in writing the initial machine development programs.

Modelling

Programs which act on the accelerator through a mathematical model of the machine are the most significant recent development in controls, pioneered at SLAC and reported elsewhere at this conference. The use of limited mathematical models has increased over the years at the SPS. From the beginning we used a simple linear model of the closed orbit, together with the measured or estimated tune value, to provide automatic orbit correction. An interesting feature here is the retention of operator verification. In practice the measured orbit and calculated corrections are displayed, and only if the results seem reasonable to the operator are the corrections applied and the resulting orbit displayed.

Right from the start the beam-line programs used a model, taking into account magnet lengths, saturation, and the beam optics, to set up the lines for different energies and different extractions during the cycle. Steering, however, was manual, and as the SPS has about a dozen beam lines this became a heavy load on the operations staff. A complete beam line model is now used, resulting in a big saving in time. The beam positions or emmittances are first measured. A model of the line is then used to optimise the steering or matching, and the corrections applied to the magnet currents. As many of our lines have over a dozen bends and quads this is a heavy computational load and takes about 50 seconds of time in FORTRAN in our NORD-10 console computers.

Another type of model, an adaptive model, is used for chromaticity correction. Here the required correction can be modelled except for three coefficients quantifying the lattice constant, the eddy current effects, and the remanent field effects. Adaption of these coefficients is done by the operator using tune measurements or observing loss effects. From the model and these three coefficients, however, the computer calculates and applies the corrections through the cycle, separating horizontal and vertical planes.

Similar partial models are applied to displays. An example is the display of bucket size, where the measured curve of RF voltage is combined with the field shape to provide a display of the variation of bucket size throughout the cycle. This is very meaningful to the operator when trying to track down losses.

Working through such partial models has three advantages. Firstly it is much faster and more deterministic than working in terms of physical quantities which can be meaningless due to coupling. Secondly it is very instructive for the operating staff who learn about the physics of the machine and do not become "knob twiddlers". Thirdly, when the models fail, a flaw in our understanding of the machine is indicated, and this can be very salutory.

A complete lattice model is outside the computational capability of the SPS control system computers. Nevertheless such a model will be required for use of the low beta insertions in storage mode operation. The approach being adopted1 is to split the model into two parts. The lattice calculation for a series of imposed low beta configurations will be calculated on the CDC 7600 using AGS. Tables of derived lattice parameters will then be transfered across data links to the SPS control computers. The on-line part of the model will then run within the control system, based on these tables, to set up the magnet currents, take measurements, and apply suitable corrections.

Surveillance and Alarms

Surveillance programs are less glamourous than the modelling or display programs, but are vital for efficient operation. Surveillance programs run more or less continuously in many of the computers and report faults to the operating staff via the alarm system. In addition they can cut off the accelerator to avoid damage or danger, for instance if a high intensity beam goes into a beam line but does not come out, or if the interlock state is incorrect.

In our experience an alarm system shows a threshhold effect in its acceptance and utility. This is determined by its credibility in the minds of the operating staff. Initially the alarm system was not good enough, often providing irrelevant or excessive information, or failing to show anything when beam loss occurred. It took substantial effort to cross the credibility threshhold and make it the essential tool it is today, not just for regular operation but also for equipment checkout and setting-up, when a close control of equipment operation is required. The main problem here is that a broad front has to be well covered, and good programs in one area can be discredited by holes in another.

The alarm system has greatly evolved over the last few years. Five lists are now maintained, Emergency, Alarms, Warnings, Experimental Areas, and a "Memo" list. An important addition has been the "Help" facility whereby a screenful (or more) of information on an alarm, eg. description, cause, action to be taken, etc. can be obtained.

Accelerator Management

Substantial effort has gone into programs which treat the accelerator as a whole. At first most of our programs were hardware oriented. We had a control system for the SPS hardware, but no control system for the SPS as an integrated whole. This "extended arm and eye" role of the control system was of course the first and necessary step in order to equal the performance of traditional control systems. There were some exceptions of course. The main power supply program just took the basic cycle parameters from the operator and automatically calculated and set up the physical parameters. Perhaps even more significantly, the beam line programs introduced the concept of file handling which has become very important.

A very considerable shock, or stimulus, to our work occurred with the advent of the thunderstorm season in mid-1977. By this time the operators had learned most of the equipment and specialists were becoming less interested in day to day running. Setting up took some time, 24 hours were allowed, but this was considered acceptable. However, a sudden thunderstorm would switch everything off, usually with some equipment damage. This episode launched a program of work which is progressing towards fully computerised setting up.

Setting-up can be divided into two parts; switch on with equipment check-out, and the loading of operational values. The latter, which we tend to call file handling, is required most often as it is used for loading machine development cycles as well as switching between operating cycles and at setting- up. A two tiered file handling system is used. Each accelerator sub-system has its own file handling for experimental purposes, tuning, etc. Then an overall system is used for the whole machine, often using floppy disks for archival storage. It is now possible to reload a complete machine from files in a few minutes, provided of course that all the equipment is operational. This results in a working machine with 60-70% transmission instead of 80%, and with beam on targets, but not of good quality. The tuning programs must then be used to get correct performance. The file handling for the individual sub-systems had been developed separately and by different people. It took a major effort by two of our operation engineers with a very broad knowledge of the accelerator to bring these together into an overall scheme. If we were to start again, stricter guidelines for file handling would be laid down from the start. Nevertheless this overall file handling system is invaluable during operation and we count it as one of the most important developments.

On the switch on front things are not so rosy. The programs are more fragmented, dealing with the hardware too much on an individual basis. Work is proceeding in two directions: firstly to provide better and more integrated sub-system switch-on programs; secondly to provide a scheduler which will bring up all systems automatically, only reverting to the individual programs when things go wrong. Another developing concept is that of the master file. This describes all the basic parameters for the required machine operating mode. All switch on and setting up programs will get their basic data from this file. Programs are provided to modify and display the contents of the file. Some programs are even capable of generating extrapolated settings using the master file for a new operating mode together with proven data from a previous mode.

The above work is difficult due to the wide range of knowledge required. We believe only engineers and operators with a wide experience can attempt it. It is also less glamourous than the tuning, measurement, and modelling, programs. Nevertheless it is of vital importance for successful operation, and is where the benefits of computer control are felt most strongly.

Programming

The work described above is essentially the performance of the programs. These can be considerd the heart of the control system. On the one hand the programs have access to the accelerator parameters via the computers, data modules, CAMAC, multiplex, etc. On the other hand we have the control centre facilities which allow the programs to communicate with people. What the control system really does for the operation, however, is determined by the quality of the programs, their coherence, and the experience of machine physics and operation built into them.

The speed of execution and ease of program writing are sensitive issues at the SPS. They are closely linked through the decision to use an interpreter. This was a fundamental design decision in the SPS control system,[2] and has proved very successful.[3] In particular, the way in which the operations staff, at all levels, have been able to take an increasing share of the programming has been very gratifying. The penalty for the advantages of the interpreter is, of course, execution speed for heavy calculations. This became a limitation as more and more modelling was introduced. The solution was to use FORTRAN for the modelling calculations. This has been adequate so far as the NORD-10 has a good FORTRAN compiler and is quite fast. FORTRAN is only used for such calculation subroutines, called from the main programs which are written in NODAL.

There is some criticism of the use of FORTRAN. We only use it for mathematical calculations where it is quite good and very fast. It is also the standard programming language for our machine physicists on the CDC 7600. In our experience the combination of NODAL and FORTRAN is easily used and powerful enough for most applications.

Waveforms and Analog Signals

A significant amount of time during operation is spent dealing with waveforms or cyclic functions of time. These range from the slow functions varying over the 12 second period of the machine cycle, down to the 2 nanosecond bunch length.

Many waveforms or functions are imposed by the control system using function generators. Designing, loading, and modifying these functions is a considerable load on the staff and the computer system. At first a general purpose hardware oriented function creation program was written, but this proved too cumbersome and slow for many applications. The pendulum then swung to the other extreme where every application, such as chromaticity correction or RF, had its own specialised function editor. Now there is a re-convergence into related groups such as acceleration, extraction, and beam lines.

Several approaches have been adopted for the viewing of such signals. Perhaps the most successful has been remote digitising, using commercial equipment, and digital transmission for display using the normal console graphics sytstem. This allows very fast signals to be observed. An interesting facility is to store a good or reference signal on the library, and to display it in red with the actual signal in green.

The digital approach is too expensive for the vast majority of our slow signals so an analog multiplexer is used. A great deal of work has gone into this system to improve its performance and make it easier to use. An important feature is the display of related signals at a single selection. A specially developed "Touch-Terminal" can select a given power supply, display digitally the required function from the computer system, then connect the analog signals at input and output of the supply to scopes. Some frequently used signals are directly wired to the control room, and the HF signals from the Faraday cage have their own two line multiplexer.

The Control Centre

The control room is housed in a separate building which also contains a reception foyer with a permanent secretary, a conference room, and a central computer terminal room. At first the control room activities were mainly technical as the consoles were used to get the machine running. As the SPS swung into full operation for physics the control room became an integrated operations centre. Initially the control room was rather sombre with dull lighting to favourise the viewing of TV screens and oscilloscopes. Reading logs, discussions, telephone calls, checking up on documentation, etc., soon became just as important as console work and many complaints were received. The control room was therefore re-decorated in warm pastel colours, and more powerful indirect lighting installed. There was a slight loss in terms of screen viewing, but this is greatly outweighed by the improved working environment.[4]

Originally the control room contained three separate consoles. Problems arose because many aspects of operation and machine development required team work. Now we have four consoles arranged in two groups of two. These form the curved ends of an oval. The straight parts are made up of the access control console and a new console for permanent displays. In the centre is a long table with space below to hold logs, documentation, cameras, cables, etc. The two curved parts can be considered as sub-centres of control. Each contains three analog signal bays and an alarm bay, flanked by two 3 bay computer consoles. This provides a powerful enough centre for a team working on a storage experiment, or doing a setting-up. Alternatively this unit can be used as two separate computer consoles plus an analog signal access point.

The computer consoles have seen little change in concept. The main improvement has been a new display system installed in 1977-78, providing full graphics and colour. A 5th console has been added in a separate room to guarantee that program development can continue even when the operational pressure blocks all four main consoles.

Acknowledgements

The work reviewed in this paper is the combined effort of the whole SPS division on the controls front. It is quite natural that the controls, particularly the programs, have seen an extensive evolution over the first few years of operation of a new accelerator with a newly conceived control system.

References

1. TRANSFERT DE DONNEES AGS ENTRE LA CDC ET LE SYSTEME DE CONTROLE DU SPS, P.E. Faugeras, CERN/SPS/AC/Int.Note/80-4

2. THE DESIGN OF THE CONTROL SYSTEM FOR THE SPS, M.C. Crowley-Milling, CERN 75-20.

3. CONTROLLING AN ACCELERATOR - THE OPERATIONS VIEWPOINT, V. Hatton and G. Shering, CERN SPS/80-12.

4. CONDITIONS DE TRAVAIL DANS LA SALLE DE CONTROLE DU SPS, J.P. Diss and E. Maquet, CERN/HS/TM/77-03.