IEEE Transactions on Nuclear Science, Vol.NS-24, No.3, June 1977

OPERATIONAL EXPERIENCE WITH THE SPS COMPUTER CONTROL SYSTEM

F. Beck, M.C. Crowley-Milling, and G. Shering
CERN, Geneva, Switzerland

The design of the SPS control system was described at this conference two years ago (see references). In 1975 the control system was assembled from its component parts into a working whole. In 1976 it was used to commission the accelerator, culminating in the last few months of regular operation for physics experiments.

A major design requirement for the control system was that it should be used for the commissioning of the accelerator. We are happy to report that this aim has been fulfilled. All systems except the RF were commissioned from the Main Control Room using the computer system. Indeed so powerful were the facilities available that much less beam time for testing was required than was scheduled, this being to the great profit of the installation program. The continuing requirement for the control system is of course to ensure the regular operation of the accelerator for experiments, and to provide sufficient flexibility to meet changing requirements, particularly for accelerator development.

In this paper we examine some of the more significant aspects of the control system. They are discussed in terms of their effect on the construction of the system, its use in the commissioning of the accelerator, and their effect on normal operation. Finally some performance statistics are given, together with an assessment of the overall control system behaviour.

Distributed Computer System

The SPS is controlled by 24 NORD-10 minicomputers, linked together by radial links to a central message transfer computer, also a NORD-10. This approach is at the opposite end of the spectrum from control by a single large central computer.

Perhaps the most important effect of this multi computer approach was the ready availability of computers early in the project for the development and testing of accelerator components. This had the indirect advantage of building up confidence and experience in computer control, particularly as the same interpretive software approach has remained constant from component testing to accelerator operation.

The use of a single type of computer everywhere has also brought advantages in hardware and software design and maintenance. The economy of using computers tailored to each application was soon found to be a false one, since machines without mass storage cannot be operated independently and machines with smaller than normal memories need special systems maintenance.

A feature of the SPS is its wide geographical distribution. The computers are also geographically distributed, with a general purpose computer in each of nine auxiliary buildings. There are additional computers in some buildings where there is a concentration of a particular type of equipment. This geographical distribution of computers was a tremendous help in the installation and commissioning of the equipment in the buildings. In regular operation, however, problems soon appeared due to the remoteness of the computers from the control centre. These are being progressively overcome by provision for remote bootstrapping to recover from software crashes, remote drum repairs to recover from parity errors, and finally a mobile spare for the worst breakdowns.

A step backward from complete geographical distribution is now being taken by the installation of a central computer with remote multiplexer links to handle access and radiation monitoring, as these services are particularly required when the accelerator is stopped and the control system could be down. In addition these links will be used to monitor the electronics of the remote computers from the centre.

Interpretive Software

One of the features of the SPS control system is the use of an interpretive language NODAL for writing the accelerator control and surveillance programs. In the design of the system it was realised that software was one of the biggest problems, particularly for a large accelerator. Its variety and complexity makes a large number of programs necessary if computer control is to realise its full potential. Hopes were pinned on the interpreter to provide the solution. Right from the beginning NODAL was available on the computers used for hardware testing as it could work quite satisfactorily on small configurations. The syntax of NODAL was chosen with a view to multi computer use. Thus when the message transfer system was commissioned, and the EXECUTE and REMIT commands added to NODAL, people moved easily and naturally from programming their individual computers to programming the complete network.

The interpretive mode of working through NODAL makes work much quicker and more satisfying even for the experienced programmer, particularly the on line editing facilities. More important, however, is that NODAL opened up programming to the hardware engineer and accelerator physicist. The control system currently uses of the order of 200 000 lines of NODAL, and there are usually several commands per line. Very little of this has been written by the professional programmers, most by people who have only ever used batch FORTRAN, and much by people who have never programmed before. About 60 to 70 people have been involved in this NODAL programming. The resulting programs show a quality and relevance which could never be matched by a professional programmer who is not intimately involved with the application area.

Another feature of the interpreter is its immediate command capability. Immediate commands are invaluable when something goes wrong and the normal programs can't cope, or aren't suitable. Then a few exploratory immediate commands can find the source of the trouble and put it right so that the programs can continue.

Distributed Data Base

NODAL applications programs access the accelerator hardware through software routines called "data-modules" located in the satellite computers. There are about 60 different data-modules in the system and any given satellite may contain around 20 of them. These data-modules provide high level access to the equipment from the NODAL program as they contain all the hardware addresses, calibration factors, limits, etc., together with the algorithms for accessing the hardware. In addition they control use of the hardware to prevent unauthorised action, put equipment out of service, etc.

The production of these data-modules was an interesting exercise. They are all we have left of traditional assembly language applications programs and come between the system software and the NODAL applications programs. It was decided to "democratise" the writing of these data-modules and about a dozen people have been involved of whom only 3 or 4 can be considered professional programmers, and all but two were very much part time. Only by providing carefully worked out standards, many support routines, and extensive personal help and computer aids for debugging, was it possible to obtain adequate results. Some data-modules are not yet up to the required standards in some respects, but nevertheless they all do their main job reliably and well.

The Main Control Room

The main control room is the show piece of the control system. Although different people have different small criticisms to make, these only underline its main success, namely that it unquestioningly accepted as the normal place to do work on the accelerator and its equipment. This is despite the fact that each of its three control consoles is only a peripheral of a computer attached to a message transfer system linked to remote computers which control the equipment.

Each console has only a few standard devices such as the touch panel, knob, track-ball, keyboard, and colour display. These come alive, however, in a large variety of ways. As the programs make extensive use of pictorial colour displays and were written by different people who are specialists in their field, a strong identification is created between the console and the system being controlled.

These facilities were excellent on their own in the commissioning of the accelerator. In regular operation, however, some deficiencies have been noted. Firstly there is a strong requirement for some permanent displays to inform the operating staff of the performance of the accelerator without their having to do active work on the consoles. These will take the form of computer updated displays and permanent display of some selected analogue signals. Secondly there is a lack of operations type programs for logging, saving and restoring overall machine conditions, and statistics. These are being written by the operations staff themselves as the needs become clearly formulated.

The Alarm System

Surveillance programs in the satellite computers check conditions regularly and report any faults to the central alarm computer which processes them and displays messages on a specialised screen on each console. This system works well and is an invaluable aid to operation.

Originally it was thought that the surveillance programs would be simple and that reduction and analysis would be done at the centre. This proved difficult, however, and would have resulted in a lot of traffic. Instead we have distributed the intelligence into the surveillance programs and use the normal NODAL message transfer facilities. Initially only hardware conditions were checked, but now many checks are being made on beam conditions such as target multiplicity, beam alignment, losses, etc.

The Message Transfer System

The message transfer system is one of the key hardware and software components of the SPS control system. Each computer is linked to the central message transfer computer by serial data links running at 750 kbits/sec. These are used to transmit 68 word "packets" of information between the computers on a memory to memory transfer basis. The packet contains 4 words of header and routing information plus up to 64 words of data. The message transfer computer uses the header to re-route the packet to the required destination. The reliability of this hardware was a major source of worry, and indeed some hardware trouble in this area initially gave considerable difficulty. We now run for several days between faults and the message transfer system is not one of our major causes of unreliability. The data-links can still give trouble, and this is aggravated by the wide geographical distribution and the implicit lack of communication for fault finding. A system of semi-automatic repeater replacement using the independent mechanism of the central computer should improve this situation.

With hindsight we can see that the control system is built in three layers. At the top is the NODAL "user" system with the interpreter and data modules. This runs on the SYNTRON II real time network operating system. This in turn runs on the hardware of the NORD-10 computers and on the data links which were made by a separate contractor. The requirement for an integrated network operating system was not recognised early in the project, but finally evolved from the packet transfer software supplied by the data link manufacturer, the computer manufacturers real time executive (modified), and a new file manager. SYNTRON II performs well, in particular having a fast response. A NODAL program in one computer can go and get a small amount of data from another and continue work in less than 40 milliseconds. For larger amounts of data the speed is limited by the software transfer rate of about 6kwords/s. As the message transfer computer can clear about 40kwords/s, six separate channels can work at full speed. The system is still undergoing development, particularly in areas such as overload recovery and diagnostic aids, but does its normal work very well.

Reliability

Hardware reliability of the computer system was not a major problem at any stage of the commissioning. Many applications programs had already been developed on a console simulation giving the impression of a multi computer system, and the software in the satellites was largely unchanged by the addition of the message transfer system. The remaining software problems were soon reduced to an acceptable level and are still decreasing.

Now that the accelerator is in full operation the computer system is taking its place as one of the larger accelerator sub-systems with its quota of hardware failures. In the first month of "real" operation, January 1977, 10 hours down time were attributed to computer trouble and 3.35 to interface faults. This was about 10% of the down time due to SPS systems and about 2.5% of scheduled beam time. The time lost per hardware fault is decreasing with experience, and much progress has been made in reducing the incidence of known weaknesses such as drum and memory errors. This trend is expected to continue. Hardware interventions on the computers are currently running at about 30 per month.

The equipment interface uses about 80 CAMAC crates and 1000 CAMAC modules for fast and special purpose applications. The main process interface, called GP-MPX, is sub-multiplexed from CAMAC and uses about 600 crates distributed throughout the auxiliary buildings. These contain about 5000 control and 1000 timing modules. Interventions on this equipment run at about a dozen per month, mainly occurring at accelerator start-up. These are equally distributed between CAMAC and GP-MPX, which suggests that the latter is about five times more reliable per unit installed.

Performance

The performance of the system as a control system depends primarily on the quality of the programs available. The performance of the computer system can be discussed in terms of response time and throughput. The initial worries were about the speed of the interpreter and the hardware link speed. With the help of strategically placed machine code routines the typical NODAL instruction time of 5 milliseconds has been quite adequate. The hardware link speed of 30kwords/s is well above what can currently be achieved due to software limitations.

The low-load response times of the system are very good. However it is easily possible to overload the system and three main bottlenecks have been identified. The first is core-load swapping in the satellites which is used to provide virtual memory. This is usually required at the start of each job and takes a simple data retrieval up from 40 to 180 milliseconds. In the one case where this problem is severe we are adding another computer to share the load. A longer term solution might be to improve the virtual memory management.

The second bottleneck is file transfer from the library. This is required when new programs are called using the touch buttons, and by programs which are repeated each accelerator cycle for updating displays. The load time of .5 seconds for a typical 1.5kword console program is quite fast enough on its own. With a lot of activity, however, delays become noticeable and repeat programs fail to run at the correct time. The solution here is to get the software transfer rate up to nearer the hardware rate as we have a factor of four in hand.

The third bottleneck is display time as all graphics are handled by the display computer and some display routines are quite time consuming. A new graphics system with a microprocessor in each display channel will help here. Also there is the possibility of putting more autonomous display capabilities into the consoles.

Conclusion

The SPS control system possesses quite a number of new and unusual features. Programmability and flexibility are its outstanding virtues, and the main price in terms of reliability is proving to be acceptable. Interest in its distinctive approach is growing at other machines, and a similar system is under construction for the CERN 28GeV proton synchrotron.

References

1. THE COMPUTER CONTROL SYSTEM FOR THE SPS, J.T.Hyman, Proceedings of the 1975 particle accelerator conference.

2. THE OPERATOR INTERFACE OF THE 300 GeV ACCELERATOR, F.Beck, Proceedings of the 1975 particle accelerator conference.

3. THE DESIGN OF THE CONTROL SYSTEM FOR THE SPS, M.C.Crowley-milling, CERN 75-20, December 1975.

4. THE NODAL SYSTEM FOR THE SPS, M.C.Crowley-Milling, J.T.Hyman, and G.C.Shering, CERN LAB II-CO/74-2, December 1974.

5. THE DESIGN AND CONSTRUCTION OF A CONTROL CENTRE FOR THE CERN SPS ACCELERATOR, F.Beck, CERN SPS-CO/76-1, April 1976.