THE CONTINUING DEVELOPMENT OF THE SPS COMPUTER CONTROL SYSTEM

P. Brummer, M. Collins, L. Jirden

Paper presented at the XIth Int. Conf. on High Energy Accelerators
CERN, Geneva, 7-11 July, 1980

THE CONTINUING DEVELOPMENT OF THE SPS COMPUTER CONTROL SYSTEM

P. Brummer, M. Collins, L. Jirden

European Organization for Nuclear Research (CERN), Geneva, Switzerland.

ABSTRACT

The SPS computer control system was commissioned in 1976, and has been continuously modified and expanded since that time, either to meet the ever-increasing needs of the accelerator and its experimental areas, or to increase the reliability of computers and interface hardware. The paper will describe the diagnostic and maintenance techniques which have been developed to allow repair and even installation to take place with little or no down-time of the accelerator. The changes planned as part of the conversion to a $p\bar{p}$ collider are described, and some statistics on installed material and reliability are given.

1. INTRODUCTION

The CERN SPS computer control system was commissioned in 1976. In the light of operating experience it has been successfully modified and expanded to meet the ever changing needs of the accelerator and its experimental areas, as well as to improve the reliability of the computers and interface hardware^{1,2}). This paper describes the techniques which allow repair, maintenance and installation to take place with little or no down-time of the accelerator. The expected changes to the control system planned as part of the conversion to a $p\bar{p}$ collider in 1981 are outlined.

2. USING THE OPERATOR INTERFACE

During the equipment commissioning phase of the SPS, the interpretive language NODAL was developed for writing a wide range of programs and testing prototype equipment. In this way the equipment designer unwittingly contributed to the design of the SPS control system. It was due to his collaboration and feedback that NODAL fairly soon became the only high level control language for the SPS accelerator. As a consequence many equipment designers were able to write their own multi-computer control programs without conventional assistance from systems applications programmers.

NODAL is provided with facilities for hardware fault diagnosis so that the operator can immediately localise a fault to a particular part of the control system. The traditional hardware/software barrier rarely exists because each operator, who is neither hardware nor software expert, rapidly acquires a system wide knowledge of the equipment under his control. Should a problem arise, then he needs only to note the ERROR or ALARM code generated by the failing NODAL program. In most cases the resulting intervention leads to the replacement of the offending unit of hardware or software. Intermittent problems require a longer period before statistical experience shows which course should be taken.

3. COMPUTER MANAGEMENT

On a daily basis, operation of the computer system is monitored by a set of NODAL utility programs which indicate the running status, down-time and restarts attributed to each

computer in the system. Operational down-time attributed to the computer system is shown in fig. 1. Frequent restarts of a particular computer may, for example, indicate the incidence of an intermittent computer or interface hardware problem or a bug in newly installed software. Such programs form part of a battery of simple structured NODAL test programs which may be called from the library disk and executed on any computer terminal in the system. Other programs aim at testing some particular aspect of the hardware or software operating environment. Particular blocking situations are logged and compared with similar situations on other computers in the system.

At the SPS the large majority of equipment has been designed to continue operation on last set values in the event of partial or total failure of the control system. This enables us to remove control temporarily while the accelerator is running. The inherent stability of the SPS has allowed an upper time limit of fifteen minutes to be arbitrarily set for operating in this degraded way. In difficult cases a spare computer is moved in and attempts are made to reproduce the fault on the new computer.

In the case of intermittent faults a strategy is devised to eliminate the fault by logical replacement of discrete parts of the hardware and software. It is important, however, not to remove an intermittently faulty computer from the operational system as it may be impossible to reproduce the fault under laboratory conditions. Switching over the computer takes less than five minutes and involves changing two connectors and reloading the software from the central library or local disk drive.

This method of investigating intermittent faults may take several days to complete but involves negligible loss of machine time: its thoroughness guarantees success.

4. ON-LINE INSTALLATION

Having developed the above techniques for on-line debugging of hardware and software faults, we found that we could use the same methods for transparent installation of new hardware or software whilst the accelerator continued to operate. Before each change the operations teams need to be confident that all equipment controlled by a specific computer is operating correctly. This involves making small but real changes to many machine operating conditions. If a problem occurs during the installation process then it would be necessary to return to the original software or hardware within a few minutes. NODAL management programs have now been written to make these procedures completely automatic so that software in remote computers can be down-line loaded and tested like any other equipment from a console in the Main Control Room. Using these techniques, as many as three computer systems per week have undergone major changes during normal accelerator operation. These unusual methods were originally contrived because of accelerator start-up problems which invariably occurred whenever new software and hardware were installed together during a machine shutdown. In most cases we have found that new software can be installed transparently during normal operation before the new hardware is commissioned in a subsequent shutdown period.

Recent changes involved installation of modern computers with semi-conductor memory to replace older computers with obsolescent drums. Such was the transparency of the change that

many users did not know that fourteen computers had been exchanged until the project had been completed.

5. MANAGING A MULTI-COMPUTER SYSTEM

For many reasons it is desirable to restrict the number of hardware and software configurations in a large system. However, individual computer systems tend to acquire personality, not only from the equipment they are driving, but also from the system programmer. With finite resources this eventually tends to limit the growth of the network. Figure 2 compares overall reliability with growth of the control system. Special systems are undesirable because each system receives personalised maintenance instead of being part of a general scheme. In this situation global design faults which would otherwise be dealt with by the computer manufacturer may remain hidden.

The early development of the SPS control system passed through such a phase as it developed from a system of individual computers into a true multicomputer control system.

During the integration phase, SPS computers were grouped rigidly into one of five hardware configuration sets.

Computer design faults are easy to detect because they occur intermittently on all computers in a particular set and have to be resolved by the computer manufacturer. Individual hardware faults occur singly and can be resolved by simple card replacement. Similarly, operating system software bugs occur on all computers loaded with a particular version of the operating system, whereas 'personality' bugs occur only on the computers driving a particular type of equipment. Machine engineers and operators using the SPS control system provide personality in the form of high level multicomputer NODAL application programs. This is not strictly true because there are still a few computer systems which contain special purpose application software written in machine code and maintained by professional programmers. Although these systems are all provided with the NODAL interpreter and linked in a standard way to the control network they cannot be managed as effectively as part of a multi-user multicomputer environment. Fortunately their number remains limited.

6. SOFTWARE QUALITY CONTROL

All software is tested in the laboratory before its installation in the operational environment. The operating system software and the interpreter are provided by a professional software team, whereas the majority of equipment interface software (data modules) is written to well defined standards by a data module programmer who is normally part of the equipment design team. This software is checked very thoroughly on a special test computer connected to the network. NODAL programs are available to test all 200 currently existing machine code functions and data modules. Initial tests ensure that the new module obeys the standards and does not destroy the operating system environment. Further tests prove that the new data module controls test equipment in the specified manner. Finally checks are made to ensure that all properties of the new software are compatible with existing NODAL application programs.

This elaborate series of tests is designed to ensure a high success rate for installing new software on an operational accelerator. At the present rate of installation, data module software can be changed once every three months. There is therefore strong pressure to build sufficient parametrization into the machine code modules so that operational changes can be made by passing parameters from the NODAL application program.

New facilities have recently been introduced to compile data modules directly from NODAL source programs. This has obvious advantages but it will not reduce the requirement for quality control and testing.

7. FUTURE DEVELOPMENTS OF THE CONTROL SYSTEM

When SPS operation recommences in mid-1981 it will require two modes of operation. The cycle by cycle mode for fixed target physics will continue but will be interspersed with proton-antiproton development periods with the machine operating in coasting mode. Such operation will require two alternative synchronisation/timing schemes. The present learning process proceeds on a cycle by cycle basis with occasional lost cycles and is adequate for an accelerator. In the case of a storage ring however there is only one 'cycle' and its loss is a serious matter. Various methods are being devised to reduce the likelihood of such loss and also to learn from failures when they occur. Equipment has been designed for cyclic digital measurement of selected analog signals. In the event of beam loss the recording of these signals will be frozen so that a post mortem analysis can be performed. During the coasting phase, control of most equipment will be inhibited by a software 'lock-up' facility at the data module interface level.

Development of the SPS as a $p\bar{p}$ collider, although not involving major reorganization or redesign of the computer network, will involve major changes at the applications level. The existing LIBRARY facilities will be doubled to allow for the large number of new NODAL application programs which will be needed for the new modes of operation. Two additional computers will be added to provide real-time control and correction of the low beta power supplies. Microprocessor based beam instrumentation systems will be installed as a natural extension of the existing control system with data module software incorporated in the microprocessor itself.

Storage ring operation implies the performance of large calculations for simulation and adjustment of operating conditions. It is intended to provide this computational facility in the form of a link with the large main frame computers, rather than by the use of non standard computers in the existing control system.

ACKNOWLEDGMENTS

The work reported here has been done by members of the Computer Controls Group with considerable help and encouragement from our colleagues in the Beam Instrumentation and Accelerator Operations Groups.

REFERENCES

- 1) M.C. Crowley-Milling, Experience with the control system for the SPS, CERN 78-09 (1978).
- 2) J. Altaber, F. Beck, M. Collins, M.C. Crowley-Milling, R. Rausch, The on-line repair, improvement and expansion of a working multiprocessor control system at the CERN 400 GeV accelerator, 3rd Int. Conf. on Trends in On-line Control Systems, Sheffield, U.K. (1979).

Fig. 1 - Evolution of accelerator down-time attributed to the control system

Fig. 2 - Reliability and growth of installed computer systems