

The New Controls Infrastructure for the SPS

P. Charrue and M. J. Clayton

ABSTRACT

A completely new control infrastructure has been installed in the SPS machine and experimental areas, replacing the old control system based on NORD computers that dated back to the 1970s. The new system uses Unix workstations and X-terminals to replace the old console computers and PC and VME chassis running LynxOS to replace the low-level interface computers.

This paper will present the old method of equipment access and then describe the transitional phase when the two systems were run in parallel followed by the final complete transition to the new system and the removal of the NORD computers.

A great effort was made to recreate the old programming environment in the new system in order to preserve the enormous investment in application programs. The equipment access and the NORD console simulator are two examples of this effort. Finally the paper will present the results of the first few months of operation of the SPS and its experimental areas with this new control infrastructure.

OLD CONTROL ARCHITECTURE

The SPS Machine

The original control system of the SPS dates from 1976, and has been extensively described elsewhere [1]. Here we will briefly summarise its main characteristics.

It was based on NORD 100 computers connected together in a star network designed specially for the project. In the control room, dedicated computers drove specialist consoles which provided a graphical user interface with trackballs and knobs, while computers installed around the machine were connected to the equipment either by CAMAC or by a fieldbus designed for the project, the Multiplex systems [MPX]. All the applications software was written in an interpretive language, NODAL [2] which had instructions that allowed commands to be executed on any machine on the network. The computers distributed around the machine, to which equipment was connected, contained special functions written in the NORD assembly language which allowed the NODAL programmer easy control of the equipment. It is interesting to note that these functions, called Data Modules, contained their own private data areas and within the limited resources of the time implemented many of the concepts now used in object-oriented languages.

Although many of the computers in the field were connected to equipment belonging to different groups and acted as general service machines, some computers were dedicated to special tasks and contained a lot of special system code. An example of this was the PS computer that controlled the magnet currents in a real time feedback loop.

The general layout of the hardware connections is shown in Figure 1. The NORD computers were connected to the CAMAC crates directly via their I/O bus, the CAMAC crates being controlled by special controllers that appeared as NORD devices on the bus. Although some equipment was controlled by CAMAC modules, most was controlled by MPX modules which were installed in crates attached to the MPX fieldbus, a serial bus with sufficient range to allow control throughout an auxiliary building.

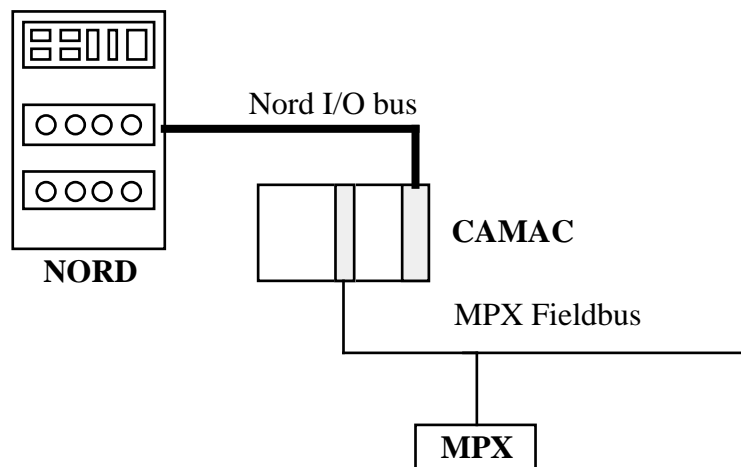


Figure 1

The SPS Experimental Areas

The SPS has two large Experimental Areas, the North and the West, and the Experimental Areas group was responsible for the control of the beamlines between the targets and the experiments. The control system was basically the same as that of the machine: it used the same NORD computers, network, programming languages and concepts. It differed in two ways from the control system of the machine: it used Serial CAMAC as the interface to the equipment and it was multi-user as it provided terminal access for each experiment. Another difference between the Experimental Areas and the Machine had a profound effect on the design and evolution of the control system and that was the fact that the Areas were constantly changing, even during operation, as experiments came and went. The system therefore contained extensive facilities for the easy installation of new equipment even during operation.

The general layout of the hardware connections is shown in Figure 2. The layout was later complicated to allow for the connection of several computers to the Serial CAMAC, but the general principle remained the same.




Figure 2

There were several Serial CAMAC loops in each zone, there being three in the largest one, the North Experimental Area. This zone also contained the longest loop which was about 2km long.

NEW CONTROL ARCHITECTURE

Although a great success, the NORD-based control system became old and difficult to maintain, and the restrictions imposed by its limited address space and non-standard operating system became more and more of a problem. The new LEP control system [3] provided a model for the integration of modern operating systems and computers in the accelerator control environment towards which the SPS control system could evolve.

The general outline of the new control architecture is shown in Figure 3. It has three levels:

- **The Control Room Layer.** This consists of X-terminals and UNIX servers which provide the Graphical User Interface to the operators, the computing power to run the application programs and specialist services like file servers, databases and alarm servers.
- **The Front End Computing Layer.** This is made up of front end process computers installed in the field and controlling clusters of local equipment. These front end computers are PCs or VME crates running a real-time UNIX (LynxOS) and tasks are assigned to them on a geographical or functional basis.
- **The Equipment Control Layer.** This layer consists of Equipment Control Assemblies (ECA) and industrial Programmable Logic Controllers (PLC) which are connected to the front end process computers by various fieldbuses or RS232 links.

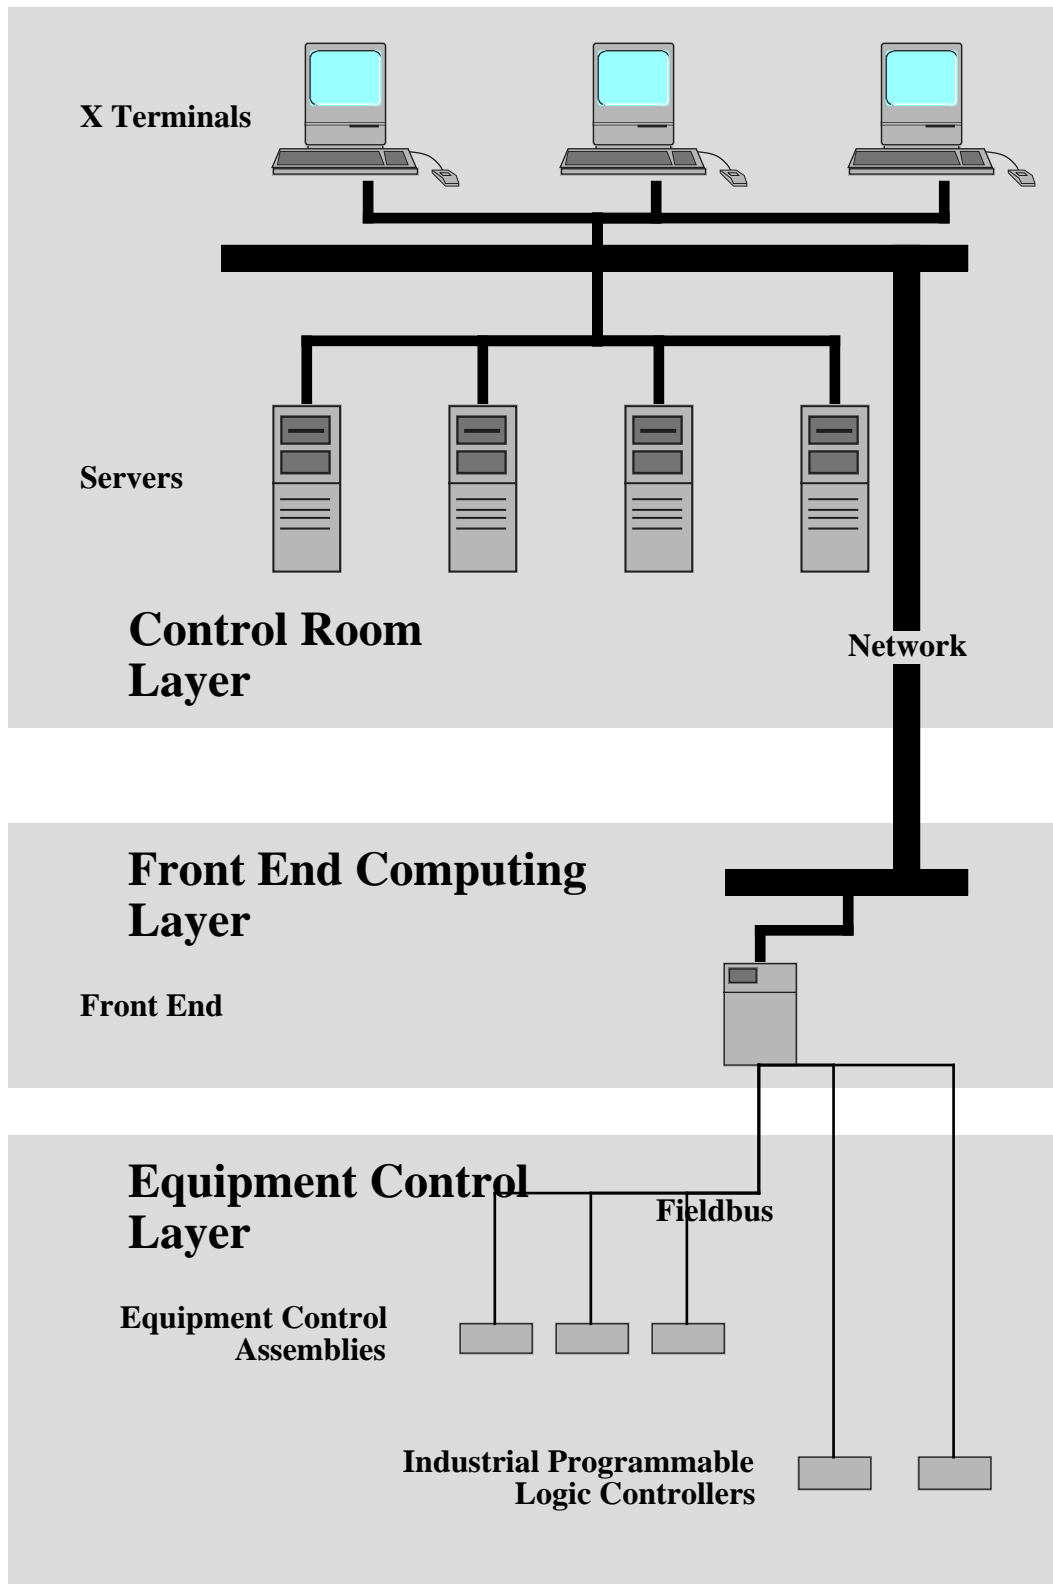


Figure 3

The network between the control room layer and the front end computing layer is a TDM/Token ring backbone. This is connected to Ethernet in the surface buildings and the Experimental Areas via IP routers and to a variety of transmission media in the control rooms via a central intelligent MMAC HUB [5] [6].

The Control Room Layer consists of X terminals, HP-UX Unix servers and some Apollo workstations. The software running at this level communicates with the Front End computers either by RPC calls to specialist servers written for the particular application, by the standard SPS/LEP EQUIP call [7] or by the network communications facilities written into the NODAL language. The SPS/LEP

EQUIP call mechanism is the most important standardisation of interprocess communications that has been introduced in the system, so a brief description is given here. An equipment database is maintained in the system, which stores all equipment by name and holds for any piece of equipment both addressing information and the method of access. A standard equipment access library is provided for C applications programmers and NODAL is in this instance another C applications program. The library translates the standard equipment access call into the particular protocol used by the actual equipment, thus removing this level of detail from the application programs. This method of access is used to communicate with the data modules in the Front End computers.

The Front End computers are 386-based PCs running LynxOS. The fieldbus uses MIL STD-1553 which is controlled by cards installed in the PC or by modules installed in the VME crates. The detailed layout is shown in Figure 4 below

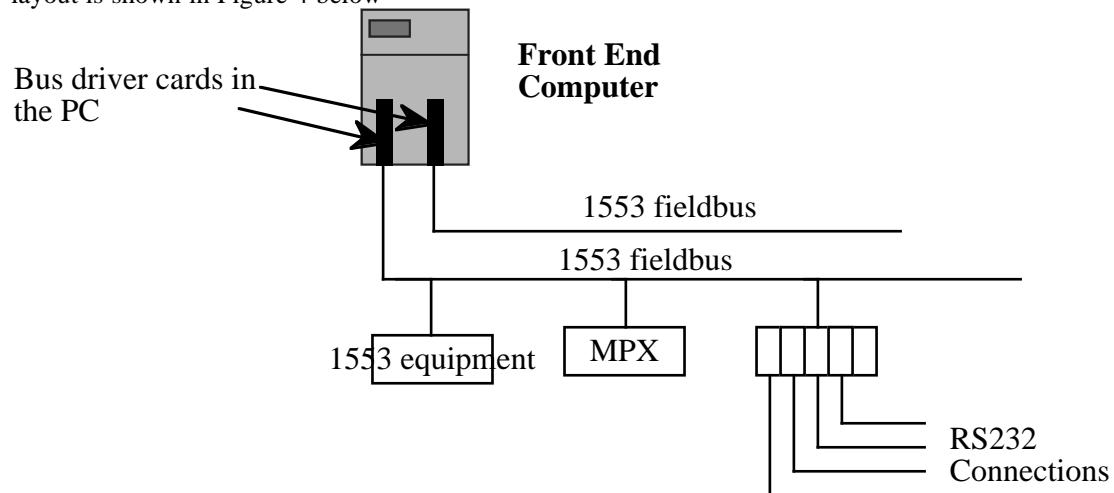


Figure 4

The software in the front end computers consists of

- Drivers for the bus driver cards.
- Translations of the original NORD data modules into C for the new environment.
- A new NODAL interpreter to execute the original NODAL code from the NORD environment.
- A message handler to route EQUIP calls.
- Any specialised equipment servers required by the local equipment.

The data modules were translated from the original NORD assembler code using a strict template which greatly speeded up the process of translation, assured a homogenous functionality of the data modules in the new system and will simplify maintenance in the future.

Some intelligent equipment had been developed for the SPS following the LEP standards, and changes had been made to the old control system to permit access to this type of equipment. Such equipment fitted naturally into the new system as it was immediately accessible via the Equip call mechanism.

The SPS Machine

The old SPS control system had many programs specifically written to take advantage of the particular environment which provided trackballs, knobs and several computer screens. The programs manipulated this environment not only by calling specialist functions but also by sending extensive sequences of control characters to the devices. Many of these programs had been written by specialists long departed and their translation to a modern environment could not be envisaged on a reasonable timescale.

The problem was solved by writing a piece of software that emulated the behaviour of the trackball and knobs and mapped the different screens of the old system into windows on a X terminal. The old NODAL programs could run as before, communicating with the operator via the emulator and with the equipment by NODAL network commands and the new data modules installed in the Front Ends.

The console emulator solved the problem of the old application programs written in NODAL. There were, and still are, a large number of application programs written in C and running on the Apollo workstations. These are the result of a first attempt to rewrite the control system in a more modern environment, but they use the proprietary graphics standards of the Apollo computer and they must be replaced by programs using the X windows graphics standards. They continue to run because they access the hardware either by calls to RPC servers that exist in the new system, or by emulating the NODAL access calls.

The original SPS machine had a very large MPX installation with about 1000 crates and at its peak over 5000 modules. Although this installation was gradually reduced as equipment groups slowly converted their equipment to be directly connected to the 1553 fieldbus (or in some cases directly to the Ethernet), a large installation still remained. This was connected to the new system by allowing MPX crates to be connected directly to the 1553 fieldbus with a new controller module, and software was written to allow access to the equipment in exactly the same way as before.

The SPS Experimental Areas

Although the old Experimental Areas control system had two notable upgrades during its lifetime, the introduction of microprocessor switch crates to share the load between several Nord computers, and the introduction of the faster Nord-100 computers in place of the old Nord-10s, it kept the basic structure of the original 1976 control system and ran with no major changes after 1980. Additions and changes were made to the large pool of applications programs.

The problem of the upgrade of the Experimental Areas was examined by a working group, which proposed the solution [4] which is now implemented. Two major problems dictated the form of the new control system for the zones: the large investment in hardware at the level of the CAMAC crates and below and the large and in some cases little understood suite of applications programs written in NODAL over the preceding 15 years by many engineers and technicians, most of whom had moved to other projects.

The old and the new control systems are shown in Figure 5.

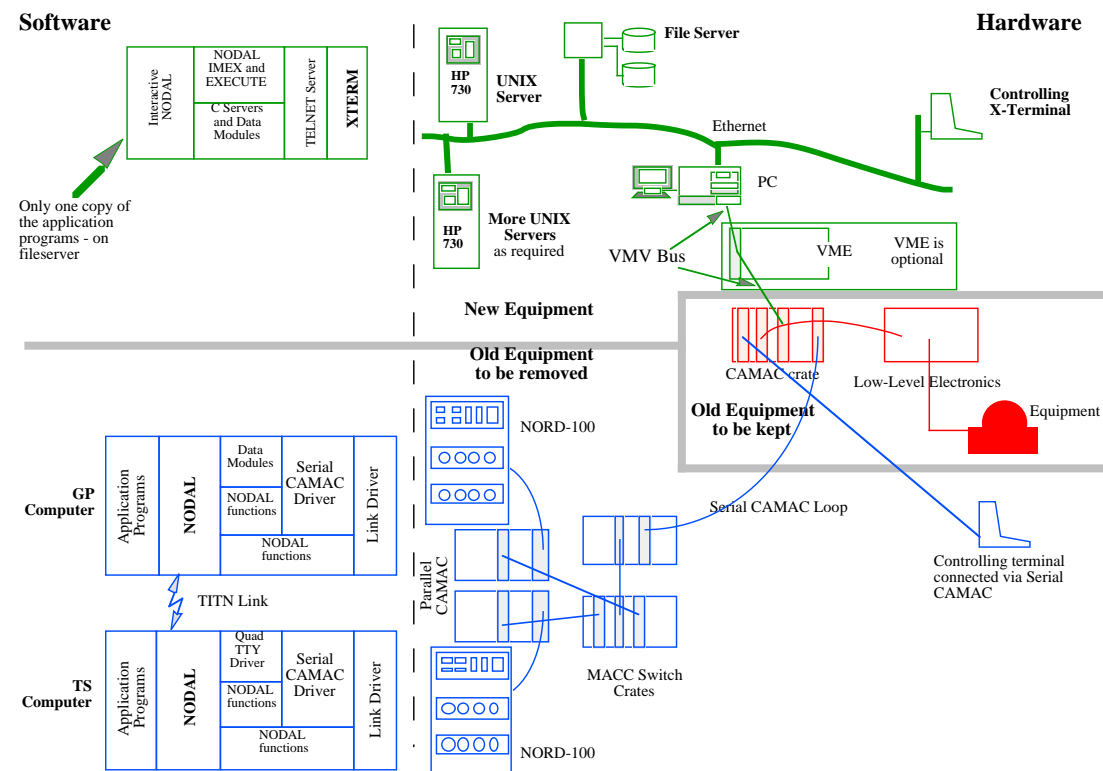


Figure 5

At the bottom of the diagram is the equipment to be removed: the Nord computers and the local parallel CAMAC, the Serial CAMAC loop and the old simple terminals. In the middle of the diagram is the equipment to be kept: the Serial CAMAC crates, the CAMAC modules installed in them and the interface electronics that was connected to these modules. This represents a very large installation as shown by the table below:

Zone	# CAMAC Crates	# NIM Crates	# Patch Panels	# Modules	# Connections
North	28	81	179	1488	8981
West	17	49	108	903	5453
EA Lab	10	29	64	531	3208

It should be noted that the EA Lab mentioned in the table was a third small installation used for test and development. It was the very large investment in material and also in its installation as shown by the large number of connections that made the preservation of this part of the installation unavoidable.

The equipment in the SPS Experimental Areas is installed large experimental halls; the largest, EHN1 is 290 metres by 50 metres, so the equipment is actually installed in small barracks which are inside the halls. This causes the equipment to be concentrated naturally into clusters called stations consisting of one or more CAMAC crates and the dependent lower level electronics. A station controls the equipment geographically close to it.

The upper part of the diagram shows the new installation. It follows the structure of the new architecture with the front-end computers in the zones and HP-UX server machines at the control room layer level. These machines are actually installed in the control room area for ease of maintenance, but they provide a terminal service all over the zones via the Ethernet. The main departure from the software structure of the machine is in the fact that the data modules run in the HP UX machines and communicate with the equipment via small specialist servers in the front end machines. This departure from the normal structure was done because it maintained the overall software structure of the old system, with a single equipment access machine for a single experimental area: an arrangement that was deeply built into the structure of the programs.

The CAMAC crates were connected to the front end PCAs by the VICbus, a parallel bus originally conceived to allow extension of the VME address space to several crates. The CAMAC crates were connected to the bus by commercially-available controllers, and it was hoped that this scheme would allow the easy mixing of CAMAC and VME and lead to a graceful future extension of the system. In practice, it has been found easier to connect VME crates directly to the Ethernet and the VICbus has not turned out to be commercially successful.

THE TRANSITIONAL PERIOD

The transitional period was a critical one, where parallel access to the equipment from the new and the old systems had to be supported as much as possible to avoid the risks and disruption of a “big bang” approach, where an old system is removed and a new system installed and started from scratch. The parallel access was successfully carried out for the machine and the Experimental Areas, although each had to have its own solution.

The SPS Machine

In the case of the machine, the original hardware layout is shown in Figure 1. This was replaced by an arrangement which allowed both the old NORDDs and the new front-end machines to be connected to the fieldbus at the same time, as shown in Figure 6.

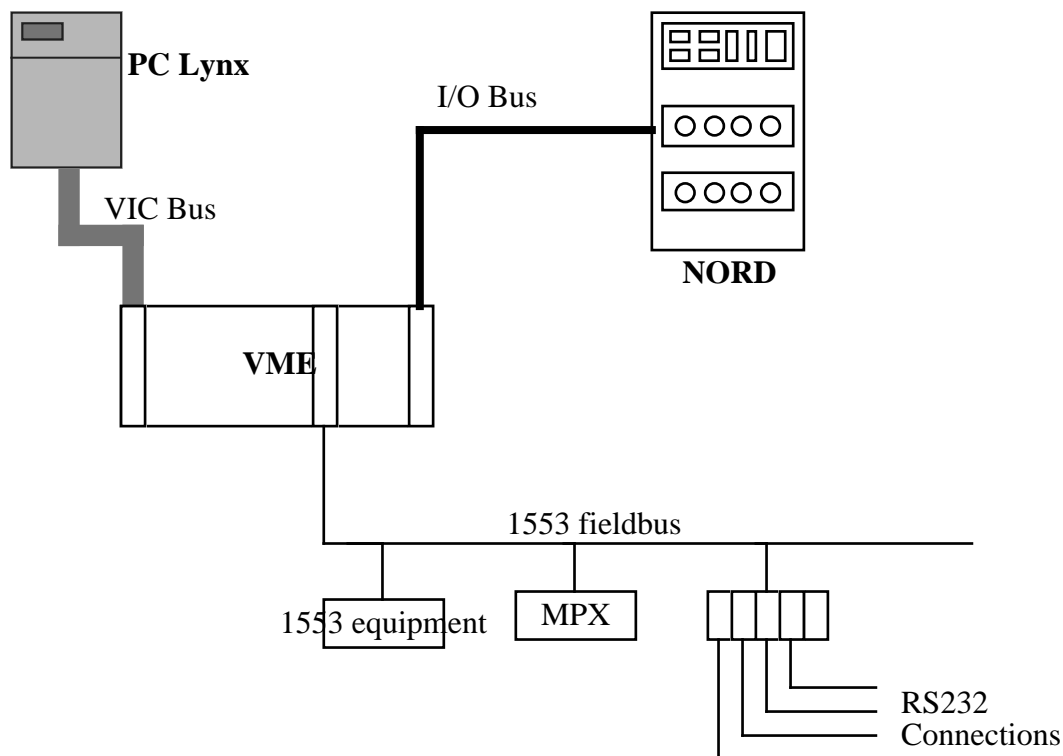


Figure 6

With this arrangement, the translated data modules could be installed and tested in the new Lynx front-end machines while the old data modules continued to operate in the NORDDs. When the new data modules became operational, changes were made in the NORDD systems to allow them to be called from

programs running in the NORDs. As the NORD systems were essentially frozen by this time, this was accomplished by patching the NORD systems so that the equipment appeared to be intelligent LEP-type equipment, for which drivers already existed in the system. The calls to this equipment were intercepted in the VME crate and routed to the new data module in the front end Lynx machine which performed the required action.

A lot of time and effort was required to tune the console emulator so that the old programs could run. This was particularly true in the case of the programs for the SPS machine as the users had over the years exploited many undocumented features of the old system and a decision had to be made each time as to whether to change the program or whether to change the emulator. Usually the decision that required the least work was taken, so that facilities that were used in many programs were reproduced in the emulator, but in some cases changes to the programs were made.

The SPS Experimental Areas

The same principle of graceful transition was applied in the Experimental Areas, but the different hardware arrangements dictated a different solution. The VICbus controller could be configured to act as an auxiliary crate controller and it was used in this position in the transitional stage as shown in Figure 7 below.

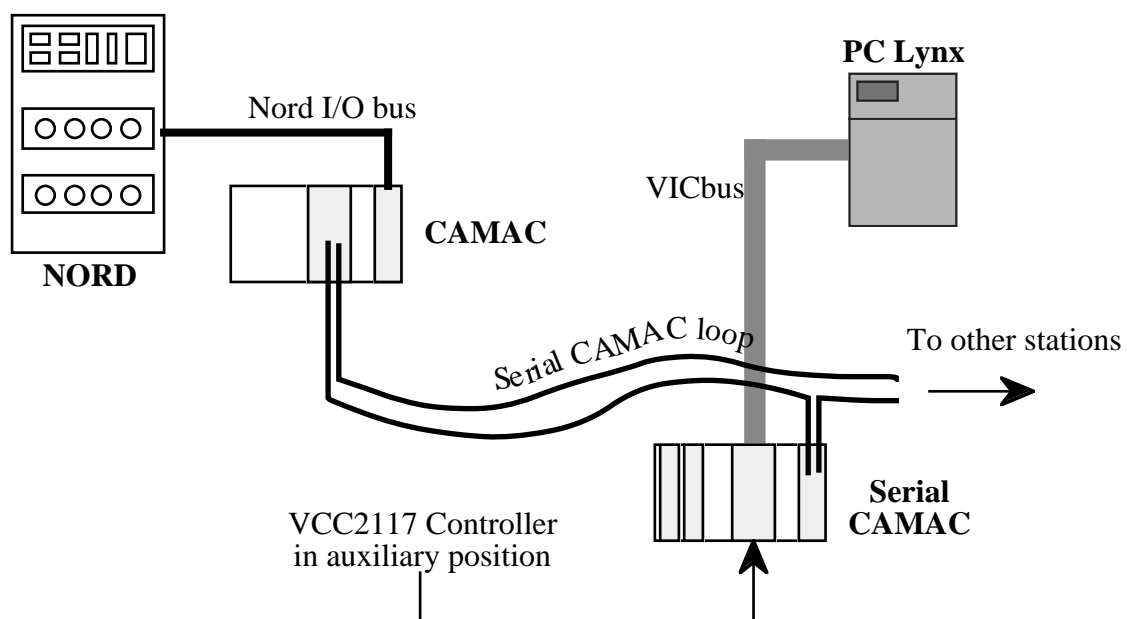


Figure 7

With this arrangement it was possible to test the new data modules while the old system continued to run. It was hoped to run the two systems completely in parallel to allow a painless transition between the old and the new systems, but this was prevented by the problem of the LAMs, the CAMAC interrupts. The CAMAC interrupt mechanism had been fully exploited in the old system, both for the connection of terminals to the system and for the operation of slow devices where the old system could not poll for the response. This mechanism was not used in the new system as the dedicated front end computers could poll without affecting the performance of the machines in the control room layer. However, when both systems were operational the interrupt was taken by whichever one saw it first, usually the NORD as it was not polling. This meant that when the old system was operating, the new system would work only partially. It could be made to work fully by stopping the old system or by disabling the interrupts in the Serial CAMAC crate controllers and it was the latter course of action that was usually followed. This disabling could be done very quickly by a program and just as quickly reversed. It was thus easy to run test sessions during machine development periods.

There were some problems with the NODAL interpreter, but fewer than those experienced with the Console Emulator as the programs were written to run on simple alphanumeric terminals and did not exploit the idiosyncrasies of the old system as much as the machine programs did.

OPERATIONAL EXPERIENCE

The new control system was put into operation for the 1994 start up of the machine, and for the 1995 start up of the Experimental Areas.

The main problem experienced during the machine start up was associated with access to the data from the target secondary emission monitors. These monitors were intelligent LEP-type equipment, and worked well during tests, but became very difficult to access during actual operation, causing programs to block while waiting for data, and finally crashing the front-end computer. The problems were traced to time-out problems in the equipment, which was unable to respond, or responded with corrupt data during a brief time after an acquisition. This problem was finally solved by queuing the requests for this equipment in the message handler and imposing a suitable delay between requests. This problem shows that no matter how extensively the equipment is tested beforehand, it is impossible to foresee the access patterns that will occur during actual operation.

The Experimental Areas, with the benefit of the experience of the commissioning of the machine, started with no real problems and beam was available for the experiments several days ahead of schedule. The remaining problems are mostly concerned with very long term stability: the correction of memory leaks and the slow accumulation of useless processes.

ACKNOWLEDGEMENTS

The authors would like to thank the SPS Operations group, the Physicists in the Experimental Areas group, the operators in the Power Converter group and all their other colleagues whose hard work in testing their programs, and in exercising the new system in both test and operational conditions was invaluable to the success of the new controls infrastructure.

REFERENCES

- [1] **The Design of the Control System for the SPS**, M. C. Crowley-Milling, CERN 75-20
- [2] **The NODAL System for the SPS**, M. C. Crowley-Milling and G. C. Shering, CERN 78-07 (1978).
- [3] **The LEP Control System**, P. G. Innocenti, Nuclear Instruments and Methods in Physics Research, ICALEPS Vancouver 1989.
- [4] **Report of the Working Group on the Upgrading of the SPS Experimental Area Controls**, H. Atherton, G. Baribaud, O. Berrig, P. Charrue, H-P Christiansen, M. J. Clayton (Chairman), W. Heinze, K. D. Lohmann, D Thomas, SL/CO/Note/91-27
- [5] **The SPS and LEP Control Network Architecture**, P. Lienard et al, Nuclear Instruments and Methods in Physics Research, ICALEPS Vancouver 1989.
- [6] **Evolution of the SPS and LEP Communication Control Network for the SPS and LEP Accelerators**, P. Lienard, Nuclear Instruments and Methods in Physics Research, ICALEPS Berlin 1994.
- [7] **The equipment access software for a distributed UNIX-based accelerator control system**, Nikolai Trofimov, Serguei Zelepoukine, Eugeny Zarkov, Pierre Charrue, Claire Gareyte, Hervé Poirier, Nuclear Instruments and Methods in Physics Research, ICALEPS Berlin 1994.