EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH

SPS DIVISION

[redacted]

CERN SPS/89-39 (ACC)

The LEP Control System

P.G. Innocenti

Abstract

The LEP Control System is based on a two level distributed architecture, with operator consoles, servers and process computers connected to the upper level and imbedded microprocessors to the lower level. A description of the system is given and performance measurements made during the first months of operation are reported. The application software used for LEP start-up is briefly described.

Invited paper for the International Conference on Accelerator and Large Experimental Physics Control Systems, Vancouver, British Columbia, Canada, October 30 - November 3, 1989.

Geneva, October 1989

The LEP Control System

P.G. Innocenti

European Organization for Nuclear Research (CERN),
CH - 1211 Geneva 23, Switzerland

The LEP Control System is based on a two level distributed architecture, with operator consoles, servers and process computers connected to the upper level and imbedded microprocessors to the lower level. A description of the system is given and performance measurements made during the first months of operation are reported. The application software used for LEP start-up is briefly described.

1. Introduction

The design of the LEP control system had to cope with constraints arising from the large circumference of the machine (27 km), its location completely underground and the small number of access shafts which can be used for exchange of control information with the surface [1]. It was possible to install only a limited number of surface links interconnecting the eight access points and the central control room; the remaining links have been established by cables laid in the tunnel.

Transmission of signals over large distances is done most efficiently by the use of optical fibres. In the LEP tunnel, however, the use of fibres manufactured according to present day technology was not possible on account of the expected damage in the presence of synchrotron radiation produced by the beams.

Finally the decision of using the existing SPS accelerator as an injector for LEP and to operate both machines from the same control room, introduces further constraints: the control systems for the two accelerators must communicate efficiently despite their marked difference and the SPS Prévessin Control Room (PCR) which sits on top of the SPS ring is far away from the LEP ring [2, 3].

2. Telecommunications Links

A cost analysis of the cables necessary to connect the LEP access points with the LEP/SPS control room (PCR) has shown that substantial savings could be obtained by installing a few high bandwidth links, either optical fibres on the surface or coaxial underground, and multiplex as many services as possible on these links [4]. Only a few tens of hardwired signals, essential for the safety of personnel, are transported over copper without multiplexing. The ultimate bandwidth for digital signals is determined by the characteristics of the installed copper or optical fibre cables between sites and PCR: it ranges between 1 and 5 Gbit/s depending on the site.

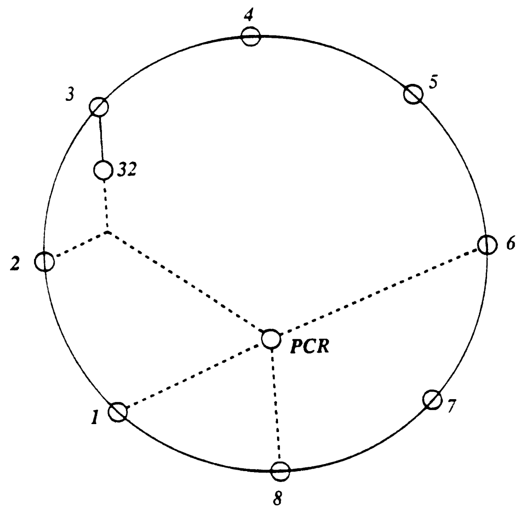

Fig. 1 High bandwidth connections by coaxial cable (solid line) and optical fibre cable (dotted line).

However, the capacity installed at present is between 68 and 140 Mbit/s, thus permitting easy extension by upgrading line equipment or adding repeaters in some cases. The method of multiplexing follows CCITT recommendations of the G700 series for Time Division Multiplexing (TDM). The TDM network carries services such as the machine control computer networks, machine synchronization, connection of experiments to the CERN computer center, terminal and voice traffic through a distributed ISPBX.

An exception to time division multiplexing is the provision of unidirectional television links from the sites to the PCR. This is done by Frequency Division Multiplexing (FDM) of frequency modulated TV signals on coaxial or optical fibre cables [5].

3. Architecture

The control system of LEP is distributed both geographically, as computers and microprocessors are installed in 35 locations, and logically, as most control actions involve many computers. In such an environment the control requirements are satisfied by connecting the computers in a network with suitable communications protocols and by synchronizing the computers by means which are independent from the network itself.

In LEP the computer network has two levels:

- The upper level which consists of central consoles and servers and local process computers, all running a program or sequence of programs under operator control.

- The lower network which consists of microprocessors imbedded in the equipment; they present fixed access points to the upper level and perform predetermined tasks, either upon request or routinely.

The synchronization system acts on computers and microprocessors by broadcasting calendar events (millisecond clock and programmed triggers) and particle bunch related triggers.

3.1 Token Passing Rings

A system of interconnected rings (IEEE 802.5) provides the support for the upper network connecting operator consoles, central servers and process computers [6]. One feature of the token ring protocol makes it particularly suitable, among other standard LAN protocols, for a large machine like LEP: one way transmission. This means an obvious but substantial saving on transmissions gear and permits easy mixing of different physical transmission media (twisted pairs, optical fibres and TDM channels) around the loop. Additional advantages are the remote monitoring built in each station and the possibility of reconfiguration by wrap-back in case of a fault, foreseen by the wiring scheme.

The rings going around LEP are approximately 40 km long. TDM channels are used in the rings for long distance transmission (1 to 5 km). Physical cable length, repeated buffering in TDM equipment and number of hosts determine a latency time of up to 400 μs in the longest ring. However, this does not penalize performance, as messages are made unavoidably long by protocols and throughput is limited by the protocols' upper layers.

Functionally, the LEP token rings are organised in three subsystems: the PCR network covering the control room; the machine network connecting hosts located underground, with functions related primarily to beam observation; and the services network connecting the surface buildings, with guaranteed high availability, as it carries information related to safety of personnel and equipment.

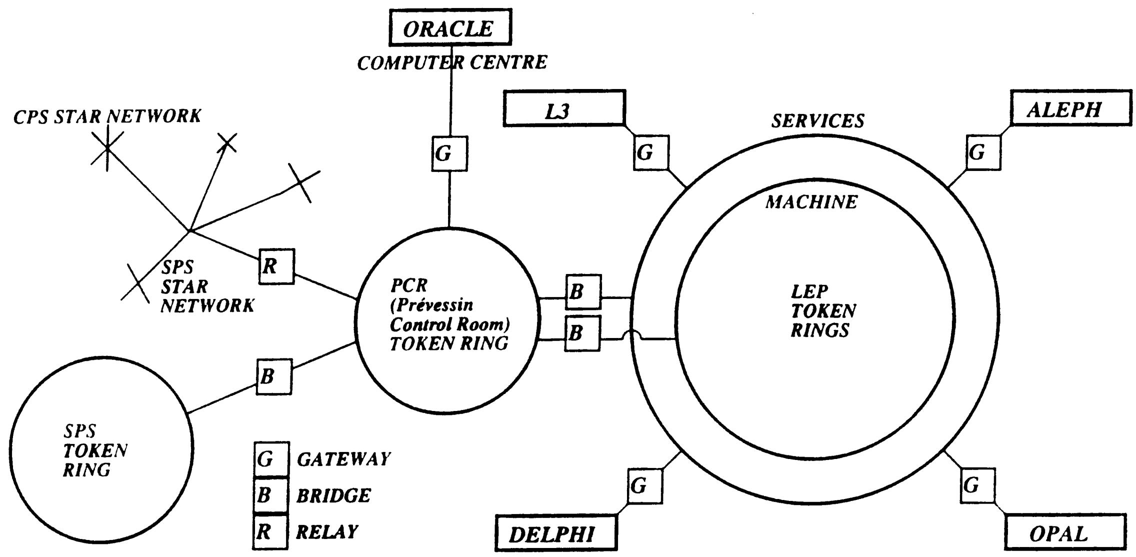

Fig. 2 LEP token-ring interconnections

Protocols from the physical level up to level 2 (LLC Class 1) conform to ISO standards and to DARPA TCP/IP for upper levels.

Facilities for Remote Procedure Calls (RPC) are available on all hosts, to permit remote execution of control programs [7].

Performance of the current version of RPCs has been measured for three situations:

- Connect time: open connection, issue one remote procedure call and close connection, 45 to 85 ms.
- Call time: issue one remote procedure call on an already open connection, 23 to 48 ms.
- Throughput: maximum transfer rate for a big buffer (40 kbytes), 25 to 37 kbytes.

The range of figure quoted also reflects the performance of the hosts, both in terms of clock speed and efficiency of the operating system [3, 6].

MAC level bridges are used to interconnect elementary rings with a throughput of 88% of the maximum ring capacity of 3.77 Mbit/s [8]. IP level gateways are used to interconnect some ring subsystems between each other or with Ethernet segments like the four LEP experiments or the computer center. Throughput of a PC/AT based gateway is 30 kbytes/s.

The number of hosts on the LEP token rings is 163. In addition, a number of machines is installed in the labs for equipment testing and program development. This makes a network of 286 hosts in total, centered around LEP. All these hosts run a version of UNIX. Thirteen bridges and 16 gateways are used for interconnections.

3.2 Multidrop highway

The connection between the process computers and the microprocessors in the field is provided by a multidrop highway (MIL-STD-1553B) [9]. A multidrop topology fits naturally the situation where many microprocessors are connected to one host in a tree structure. The particular multidrop highway chosen has features which make it suitable for the LEP environment:

- noise immunity, by transformer coupling,
- speed, 1 Mbit/s up to 400 m,
- distance capability, up to 20 km with repeaters, at 125 bit/s,
- single twisted pair cable,
- polling by a single master, ensuring positive monitoring.

The software protocol on the multidrop highway has been designed to be simple command-response, but has grown more elaborate to provide more services to the microprocessors in the field [10].

End-to-end communication on the multidrop highway depends strongly on the power of the microprocessors connected and on the number of drops. Typical times for an elementary transfer of 25 bytes is 80 ms and throughput on a single highway is 22 kbytes/s, in command-response mode.

4. Computers

4.1 Process Computers

The 56 process computers bridging across token rings and multidrop highways are Olivetti 380C PCs, running XENIX SCO 2.3 [11]. They are driving a VMEbus crate to which up to 8 multidrop highways are connected. The link between the PC and VMEbus is serial, using the same MIL-STD-1553B protocol. All the multidrop highway drivers (bus controllers) are based on a Motorola 68010 microprocessor and perform autonomous polling and monitoring of the stations. Programming languages are C and NODAL [12], the interpreter on which the SPS control system is based.

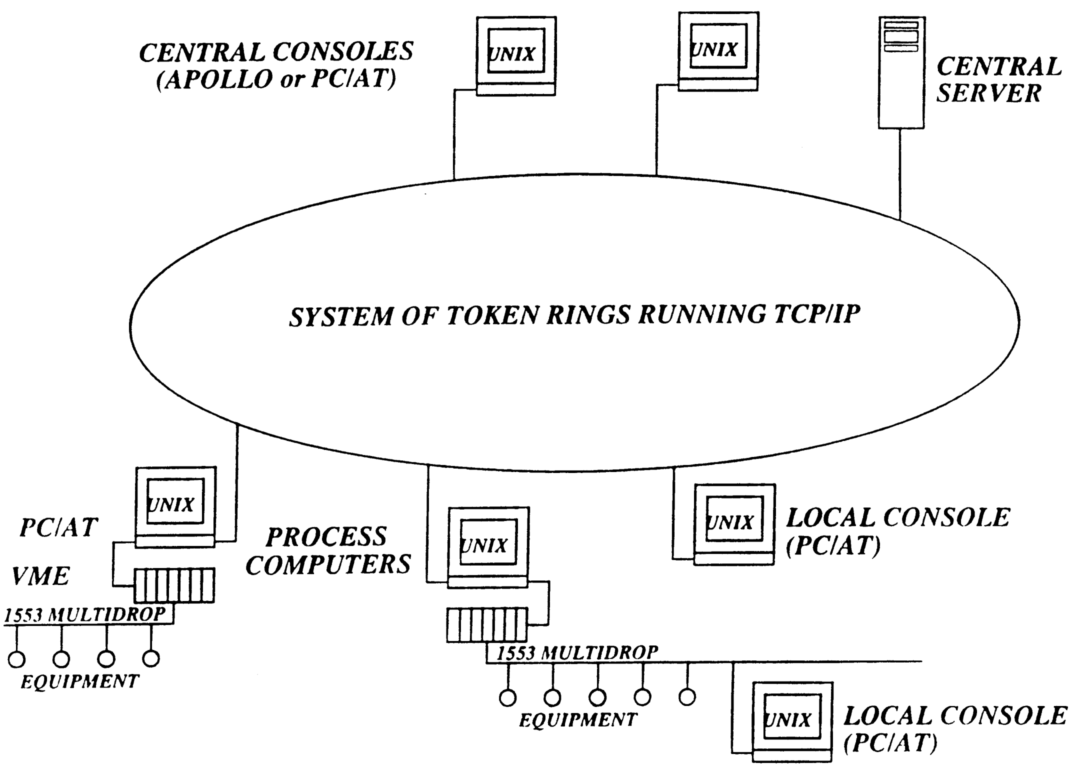

Fig. 3 Consoles, servers and process computers interconnections

4.2 Operator and Maintenance Consoles

A design goal of the LEP control system is to provide consoles which are upwards compatible from the field to the control room, so that programs which run from (simple) consoles in the vicinity of the equipment may as well run from the control room [13]. The design goal is a subset, imposed by economy, of the requirement that every program can be run from any console anywhere, where the console is a workstation. Two types of consoles have been retained: affordable PC clones [11] connected either to the token ring or to the multidrop highway, running XENIX SCO, or Apollo 3000 or 3500 (for the control room only) connected to a token ring, running UNIX BSD 4.2 or System V.3. Despite the common UNIX environment, two problems must be solved to achieve the design goal.

a) The network protocol on the multidrop highway is not based on TCP/IP. Hence for a console on the multidrop to perform services implying RPC, one must bridge the gap between the simple multidrop protocol and the RPC facilities. This has been achieved by relying on a server in the process computer.

b) Standardization of graphic packages has not materialized during LEP construction and the use of proprietary graphic presentation packages makes program portability between Apollo UNIX and PC XENIX problematic: all hopes are now placed on X Window.

4.3 Central Servers

The LEP control system does not use super-minicomputers or mainframes as central servers, but again a number of PCs or workstations. This is the case for generating synchronization patterns, for centralising beam observation and for computer network management where a number of PCs, up to the top of the range, running XENIX SCO [11] have been used; for file, archive, alarm and local data base servers with IBM 6150s, running AIX, and for file servers with Apollos.

5. Microprocessors in the field

All equipment in the field is controlled by an imbedded microprocessor, housed in general in a VMEbus or G.64 crate (~ 2000 units). In the 8-bit range, the preferred type is the M6809 (mostly running AMX), but the TMS9950 and Z80 are also present in numbers. The 16/32 bit range is exclusively covered by the Motorola 680x0 family, with RMS68k or OS.9 as operating systems. The equipment to be controlled varies in complexity from simple arrays of input/output to elaborate networks, with distributed intelligence, performing a sequence of operations or autonomous surveillance [9].

Special problems arise from commercial equipment supplied with a turn-key control system: an interface to the 1553 multidrop is needed, usually through a protocol converter.

6. Machine Synchronisation

There are two distribution systems for machine synchronization frames. One is providing General Machine Timing (GMT) consisting of pulses at one millisecond intervals (with a jitter smaller than .5 μs) and coded events inbetween [14]. It is broadcast to all hosts (consoles, servers and process computers) and to all microprocessor controllers of equipment which may need it. Long distance distribution of GMT is carried by a TDM channel, whilst local distribution is done either directly or as a companion to the 1553 multidrop, on a separate twisted pair of the multidrop cable.

The second distribution system (Beam Synchronous Timing (BST)) permits tagging each of the four electron and four positron bunches, within a revolution lasting 89 μs [15]. It is transported around the ring on a TDM channel with pick-offs in 24 locations and returned to the source.

7. Application Software

7.1 Databases

The LEP project has made extensive use of the ORACLE relational data-base management system for describing lattice, geometry, equipment para-meters and for keeping track of installation and planning [13]. All the information necessary to run the machine is stored in the database. As the control system acts on specific systems, databases must be built for each system by combining the available data; in particular a dictionary is essential to establish a relation between the various naming conventions. During operation a snapshot of beam and machine parameters is taken, either for permitting further offline studies or to be used as input to further runs. These data are saved in the control system and an archive index is stored in ORACLE: retrieval of a specific situation will be eased by the relational nature of the database management system. ORACLE is not accessible online by the control system as it runs on the Vax cluster of the CERN computer center. Therefore a subset of the database is loaded into the control system at the beginning of each run and data are accessed online from local files. Conversely, data as the archive index or the alarm history are initially stored locally and transferred to ORACLE at regular intervals.

7.2 Programs for Operation

Broadly speaking, there are two programming environments: one for the imbedded microprocessors near the equipment and a general UNIX environment for all the computers.

Programs for the microprocessors have been written primarily by the equipment specialists and both native development system and cross software development tools have been used. Most programs are written in Pascal; the executable code is either stored in local PROMs or downline loaded from a server.

The writing of most programs for the UNIX environment has been preceded by a detailed functional study using SASD methods. This has favoured a fast and efficient writing of the code and a realistic planning.

A first layer of application programs consist of servers running in the process computers. These servers are called by RPCs either from the algorithmic part of the application or directly from the presentation layer [17]. Presentation of the machine data is done by using Dataviewer [18], a package designed to display accelerator parameters.

The programs for controlling the service infrastructure of the machine have been written by relying heavily on a presentation package for industrial equipment, DV-draw, which permits easy interaction with the equipment directly from the screen of a workstation.

A global error reporting system has been implemented to log and report errors arising from the execution of control programs [17]. This is complementary to the alarm system which informs the operator of alarms raised by surveillance programs of a diverse nature around the machine [19].

References

1. G. Bachy et al., Proc. of the 14th Int. Conf. on High-Energy Accelerators, Tsukuba 1989, Gordon and Breach (to be published).

   LEP Design Report, Vol.1: The LEP Injector Chain, CERN-LEP/TH/83-29, CERN/PS/DL/83-31, CERN/SPS/83-26, LAL/RT/83-09 (1983); Vol. 2: The LEP Main Ring, CERN-LEP/84-01 (1984) (unpublished).

2. M.C. Crowley-Milling, IEEE Trans. on Nucl. Sci., Vol. NS-30, 4, p. 2142-2146 (1983).

   J. Altaber, Proc. of the 13th Int. Conf. on High-Energy Acc., Novosibirsk, Publishing House Nauka, Siberian Div., Vol. 2, p. 207-212 (1987).

   C.G. Beetham et al., Proc. of the 1987 IEEE Part. Acc. Conf., Washington, D.C., Vol. 2, p. 741-743 (1987).

   P.G: Innocenti, Proc. of the EPS Conf., Villars/Ollon (1987) (to be published).

   P.G. Innocenti, Proc. of the Int. Conf. on the Impact of Digital Microelectronics and Microprocessors on Particle Physics, Trieste, World Scientific, Singapore, p. 73-80 (1988).

3. P.G. Innocenti, Proc. of the 14th Int. Conf. on High-Energy Acc., Tsukuba 1989 Gordon and Breach (to be published).

4. C.R.C.B. Parker, Proc. of the EPS Conf., Villars/Ollon (1987) (to be published)
   C.R.C.B. Paker, Proc. of the 1989 Part. Acc. Conf., Chicago (to be published).

5. A.K. Barlow, Proc. of the 3rd Symposium on Electronics in Communications, Porto, INESC, sec. C.1, p. 1-4.

6. R. Rausch, Proc. of the EPS Conf., Villars/Ollon (1987) (to be published).

P. Liénard et al., Proceedings of this Conference.

7. P. Anderssen et al., Proceedings of this Conference.

8. M. Mignacco, Proc. of the 3rd Symposium on Electronics in Communications, Porto, INESC, sec. C.10, p. 10-16.

9. R. Rausch, Proceedings of the Conference on MIL-STD-1553B and the Next Generation, London (1989) (to be published).

10. A. Bland, Proc. of the EPS Conf., Villars/Ollon (1987) (to be published).

11. L. Jirdén, Proceedings of this Conference.

12. M.C. Crowley-Milling and G. Shering, CERN 78-07 (1978).

P. van der Stok, CERN-LEP Controls Note 63 (1986) (unpublished).

13. P.S. Anderssen, Proc. of the EPS Conf., VIllars/Ollon (1987) (to be published).

14. C.G. Beetham et al., Proc. of the 1987 IEEE Part. Acc. Conf., Washington, D.C., Vol. 2, p. 766-768 (1987).

15. C. Bovet, Proc. of the 13th Int. Conf. on High-Energy Acc., Novosibirsk, Publishing House Nauka, Siberian Div., Vol. 2, p. 220-228 (1987).

J. Borer et al., Proc. of the 1987 IEEE Part. Acc. Conf., Washington, D.C. Vol. 2, p. 778-782 (1987).

G. Baribaud et al., Proceedings of this Conference.

16. J. Poole, Proc. of the 4th European Oracle User's Group Conference, Gleansdale, UK (1986).

17. R. Keyser et al, Proceedings of this Conference.

18. D. Brignon, CERN/SPS/ACC/Note 88-1 (1988) (unpublished).

A. Sweeney, CERN-LEP Controls Note 107 (1989) (unpublished).

19. R. Nagel and M.W. Tyrrel, CERN-LEP Controls Note 89 (1988) (unpublished).