Particle Accelerators, 1990, Vol. 29, pp. 183–190
Reprints available directly from the publisher
Photocopying permitted by license only

THE LEP CONTROL SYSTEM: ARCHITECTURE, FEATURES AND PERFORMANCE

P. G. INNOCENTI
CERN - European Organisation for Nuclear Research - CH-1211 Geneva 23 - Switzerland

Abstract: The architecture of the LEP Control System is strongly influenced by the large size of the accelerator and by the (economy) requirement of multiplexing many signals over a limited number of high bandwidth links. The use of microprocessors as local controllers of practically every piece of equipment offers unprecedented distributed control and surveillance possibilities, but at the same time enhances problems of communications and date coherence. Performance of the system is analysed in the light of the design goals and of experience gained during the LEP start-up.

1. INTRODUCTION

There are two classes of constraints imposed to the LEP Control System. A first class arises from the geometry of LEP expressed by its large circumference of 27 km, by its location completely underground, with eight access shafts suitable for exchange of control information between surface and underground [1,2]. The limited number of surface links of any type between the eight LEP sites (some sites can be reached only via the tunnel) and the impossibility of using optical fibre cables in the machine tunnel on account of damage due to synchrotron radiation complete the physical environmental constraints.

The second class of constraints comes from the design choice of using the existing SPS accelerator as an injector to LEP and to operate both machines from the same control room, the existing SPS control room [3,4,5,6].

2. COMMUNICATION INFRASTRUCTURE

The prohibitive cost of control cables over distances of 4 to 10 km, as required to connect the LEP access points with each other and with the LEP/SPS control room (PCR), has strongly influenced the choice of multiplexing a large number of signals over a limited number of high bandwidth connections, either fibres on the surface or coaxial underground [7]. Only a few tens of hardwired signals, essential for the safety of personnel, are transported over copper without multiplexing.

The ultimate bandwidth permitted by the installed copper or optical fibre cables between PCR and the sites is between 1 to 5 Gbit/s depending on the site. The method of multiplexing follows CCITT recommendations of the G700 series for Time Division Multiplexing (TDM). The TDM system carries a number of services such as machine control networks, machine synchronisation, connection of experiments to the CERN computer centre, voice and terminal traffic through a distributed ISPBX.

3. CONTROL SYSTEM ARCHITECTURE

The functional requirements of controlling an accelerator in a distributed environment as in LEP are satisfied, in general, by a network of computers and by a synchronisation system.

In LEP, the computer network has two logical levels:

- the upper level consists of central consoles and servers and local process computers, all running a program or a sequence of programs under operator control;
- the lower level consists of microprocessors imbedded in the equipment, presenting fixed access points to the upper level and performing predetermined tasks, either upon request or routinely.

The synchronisation system, acting on computers and microprocessors, broadcasts calendar events (millisecond clock and programmed triggers) and bunch related triggers.

3.1. Computer Networks

a) Token Passing Rings

A system of interconnected rings (IEEE 802.5) provides the support for the upper network interconnecting operator consoles, central servers and process computers[8]. 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 savings on transmissions gear and permits easy mixing of different physical transmissions 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.

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

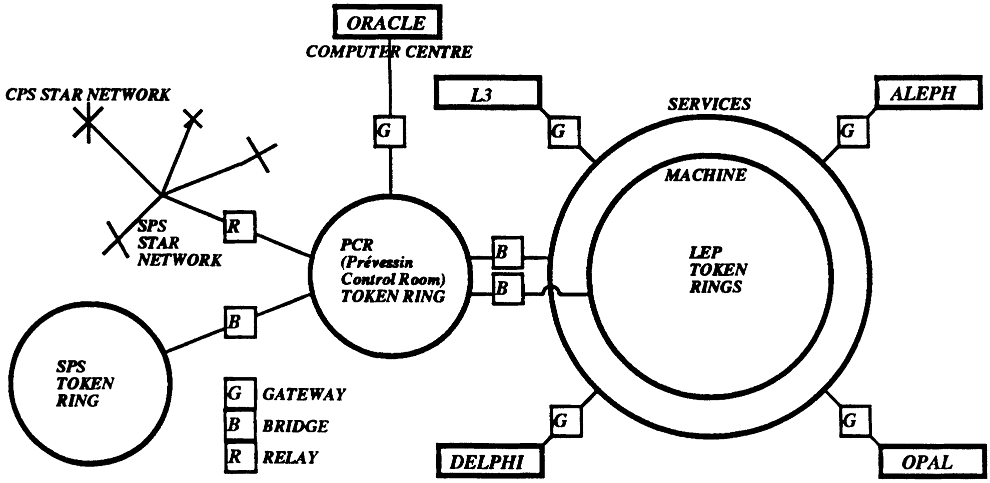

Fig. 1 LEP token-ring interconnections

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. The RPC is interfaced to UDP.

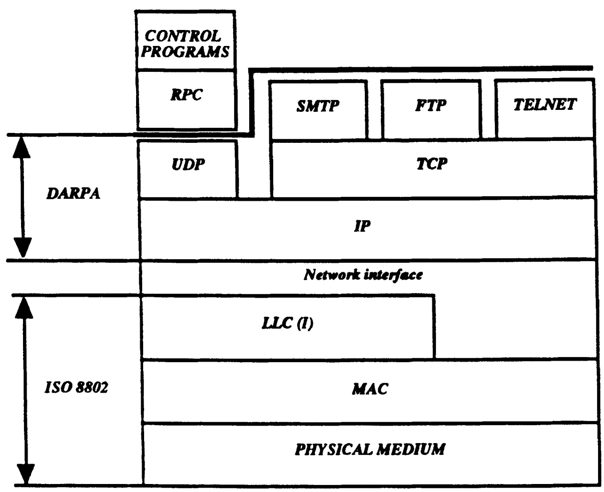

Fig. 2 Protocol organization

When using RPCs, establishing a connection between two hosts, sending an elementary message of 80 bytes and releasing the connection requires 85 to 600 ms; an elementary message of 80 bytes, once the connection is established, is transferred in 25 to 120 ms. Program-to-program transfer rates for 2 kbytes long messages are 5 to 30 kbytes/s. The range of figures quoted reflects the performance of the hosts between which the communication is established both in terms of clock speed (6 to 25 MHz) and efficiency of the operating system. Performance has been measured during LEP commissioning without any software tuning and is expected to improve rapidly.

MAC level bridges are used to interconnect elementary rings with a throughput of 85 % of the maximum ring capacity of 3.7 Mbit/s. 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 gateways 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. 12 bridges and 16 gateways are used for interconnections.

b) Multidrop Highway

The connection between the process computers and the microprocessors in the field is provided by a multidrop highway (MIL-STD-1553B). 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 LEP:

- 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 a deterministic response and 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 [9]. It is regrettable that LEP could not take advantage of the standardisation efforts of the MAP community for a field bus and protocol.

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.

3.2. Process Computers

The 56 process computers bridging across token rings and multidrop highways are Olivetti 380C PCs, running XENIX SCO 2.3. 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, the interpreter on which the SPS control system is based [10,11].

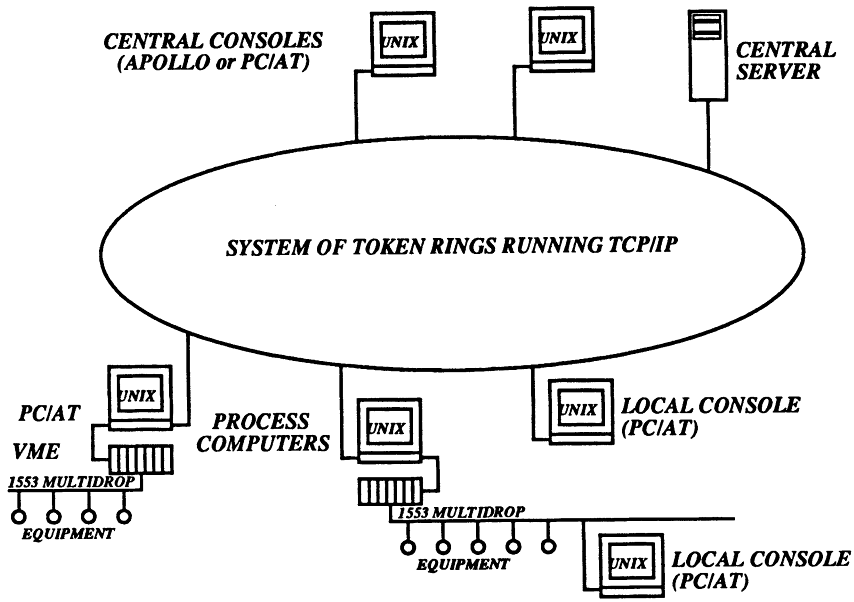

Fig. 3 Consoles, servers and process computers interconnections

3.3. 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.

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

3.4. Operator Interface

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. This 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 [12]. Two types of consoles have been retained: affordable PC clones connected to either the token ring or to the multidrop highway, running XENIX SCO, or Apollos 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 computers.

b) Standardisation of graphic packages has been slower than LEP construction and proprietary graphic presentation packages make program portability between Apollo UNIX and PC XENIX problematic. Only now, after LEP completion, it appears that it will be possible to move to a standard, X-window.

3.5. Central servers

The LEP control system does not use super-minicomputers or mainframes as central servers, but again a number of PCs or workstations. For generating synchronisation patterns, for centralising beam observation and for computer network management a number of PCs, up to the top of the range, running XENIX SCO; for file, archive, alarm and local data base servers IBM 6150s, running AIX; for file servers also Apollos.

3.6. Machine Synchronisation

There are two distribution systems for machine synchronisation 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 [13]. It is broadcast to all hosts (consoles, servers and process computers) and to all micropocessor 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 Synchronisation Timing (BST)) permits tagging each of the four electron and four positron bunches, within a revolution lasting 89 μs. It is transported around the ring on a TDM channel with pick-offs in 24 locations and return to the source. It is used primarily by the beam observation stations which are storing data in a cyclic memory in a free running mode and insert time markers to correlate the measurements with the particle bunches [14,15]. The ability of recording the position of all the bunches for many subsequent turns at each transit through the 504 position pick-ups is a unique diagnostic tool of LEP. It is made possible by the use of flash ADCs. Signals from the pick-ups are amplified and digitised in less than 450ns so that resolution in time of bunches of opposite charge is possible by this method (narrow band processing) except in the immediate vicinity of the interaction regions. Bunch tagging by BST is used also for the bunch intensity measurements by transformers.

4. APPLICATION SOFTWARE

4.1. Data Bases

The LEP project has made extensive use of a relational data base management system, ORACLE, for describing lattice, geometry, equipment parameters and for keeping track of installation and planning. All the information necessary to run the machine is stored in the database. As the control system acts on specific systems, databases by system must be built 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 on line by the control system as it runs on the Vax cluster of the CERN computer centre. Therefore a subset of the database is loaded into the control system at the beginning of each run and data are accessed on line from local files. Conversely, data as the archive index or the alarm history are initially stored locally and transferred to ORACLE at regular intervals.

4.2. 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 specialist and both native development systems 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 of the software project. 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 interaction with the equipment directly from the screen of a workstation.

REFERENCES

1. G. Bachy et al., in Proc. of this Conference (1989)

2. LEP Design Report, Vol. 1: The LEP Injector Chain, CERN-LEP/-TH83-29, 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) (unpubl.).

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

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

5. P. G. Innocenti, in Proc. of the EPS Conf., Villars/Ollon (1987) (to be publ.)

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

7. C.R.C.B.Parker, in Proc. of the 1989 Part. Acc. Conf., Chicago-Ill.(to be publ.)

8. R. Rausch, in Proc. of the EPS Conf., Villars/Ollon, (1987) (to be publ.)

9. A. Bland, in Proc. of the EPS Conf., Villars/Ollon, (1987) (to be publ.)

10. M. C. Crowley-Milling & G.Shering, CERN 78-07 (1978)

11 P. D. V. van der Stok, LEP Controls Note 63 (1986) (unpubl.)

12. P. S. Anderssen, in Proc. of the EPS Conf., Villars/Ollon,(1987) (to be publ.)

13. C. G. Beetham et al., in Proc. of the 1987 Part. Acc. Conf., Washington DC, Vol.1 p.766 (1987)

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

15. J. Borer et al. in Proc. of the 1987 Part. Acc. Conf., Washington DC, Vol.2 p.778 (1987)