SPS and LEP Controls, Status and Evolution towards the LHC Era.

R. J. Lauckner and R. Rausch, CERN, Geneva, Switzerland

ABSTRACT

LEP was the first CERN accelerator to be built with a Control System based largely on off-the-shelf technology. Operational experience has confirmed the strategies of the original design. However, to maintain this system components have to be renewed, and in order to meet the increasing user demands the system requires continuous investment. These factors have ensured that the Control System is still based on a modern computer infrastructure, and the recent extension to the SPS machine and its Experimental Areas has been achieved without any specific developments of the basic infrastructure. Increasingly, complete industrial systems are replacing old and obsolete equipment of the accelerators and technical services. Strategies are evolving to interface and progressively integrate Industrial Control Systems into the existing infrastructure. We present a status report of the Control System today and describe the current work to upgrade and prepare it for future applications in the LHC era.

1 INTRODUCTION

In July 1989, LEP was commissioned with a Control System based largely on off-the-shelf technology [1]. Communications were centred on 3 Token Rings interconnected by bridges with a maximum throughput of 3.77 Mbit/s. This system served to interconnect 163 operational host machines. Gateways to other networks such as the computer centre and the LEP Experiments were PC/AT based with a throughput of 30 kbytes/s. There were 56 front-end process computers consisting of 386 PCs linked to VME crates containing the MIL-1553-B fieldbus drivers for the connection of equipment. Consoles, all running UNIX, were Apollo 3000 or 3500 running BSD 4.2 or Sys V.3 in the control room and PC clones running SCO XENIX 386 in the remote buildings. The Apollo machines had an underlying Aegis operating system. At that time, there was no recognised graphics standard available. Apollo graphics were based on powerful but proprietary tools which prevented the use of that software in the XENIX PCs. Although mainly off-the-shelf, the Control System was designed and assembled by CERN engineers and technicians. Industrial systems were employed for some of the technical services but, again, integration work was undertaken at CERN.

While the overall architecture of the Control System was validated in 1991 [2] during the design of the Control System upgrade for the SPS and remains intact today, the components have changed rapidly. In the subsequent sections and for each item we will describe the evolution to the current state, our operational experience and the related future developments to prepare the SPS and LEP accelerator Control Systems to match the LHC requirements.

2 NETWORK

A brief description of the original LEP control network is given above and further details may be found in the literature [3]. A considerable evolution has occurred in the past 5 years. The single bridged network has evolved to an accelerator internet with routers replacing the bridges. This gives better isolation and, thus, more traffic can be handled. Following the limited commercial success of Token Ring, Ethernet was introduced using Token-Ethernet Bridges in order to connect VME-based OS9 systems for a second generation LEP closed-orbit monitoring system. Later, hubs were added to the Prevessin Control Room (PCR) to connect equipment without a Token Ring interface. The Control Rooms are now equipped with category 5 cabling, each peripheral has a dedicated cable to the hub which sits on the backbone. This step gave an important improvement in reliability and bandwidth and opened the upgrade path to new technologies such as FDDI. More details on this work were reported at ICALEPCS'93 [4].

The PCR today has a 100 Mbit/s FDDI backbone, and the main Oracle servers on the Swiss side of the CERN site are connected directly with FDDI links. A major extension to the PCR will be carried out this winter to centralise the secondary beam control for the SPS. The associated new servers and X Terminals will be connected with 100VG-AnyLAN. All these improvements are being driven by an increasing demand for network resources. The total number of hosts connected to the SL Control Network is now around 1200 and the use of the system becomes more intense as computing becomes more distributed. The first phase of the installation of an FDDI backbone for the LEP machine network was completed this spring [5]. A novel feature is the use of the drains of the LEP tunnel to carry the radiation-sensitive fibres to the alcoves of LEP. We have chosen switching routers for this installation to allow for future extensions towards virtual networking and multi-protocol networking. The system is being equipped with a central concentrator and 8 hubs switching between the radial FDDI lobes and azimuthal Ethernet in the tunnel drains as shown in Figure 1. This improvement provides the 400 hosts around the LEP ring with a shared bandwidth of 100 Mbits/s compared with. 4 Mbits/s at present. It also opens the way for the introduction of ATM technology as our requirements and the solutions available on the market continue to develop in future years.

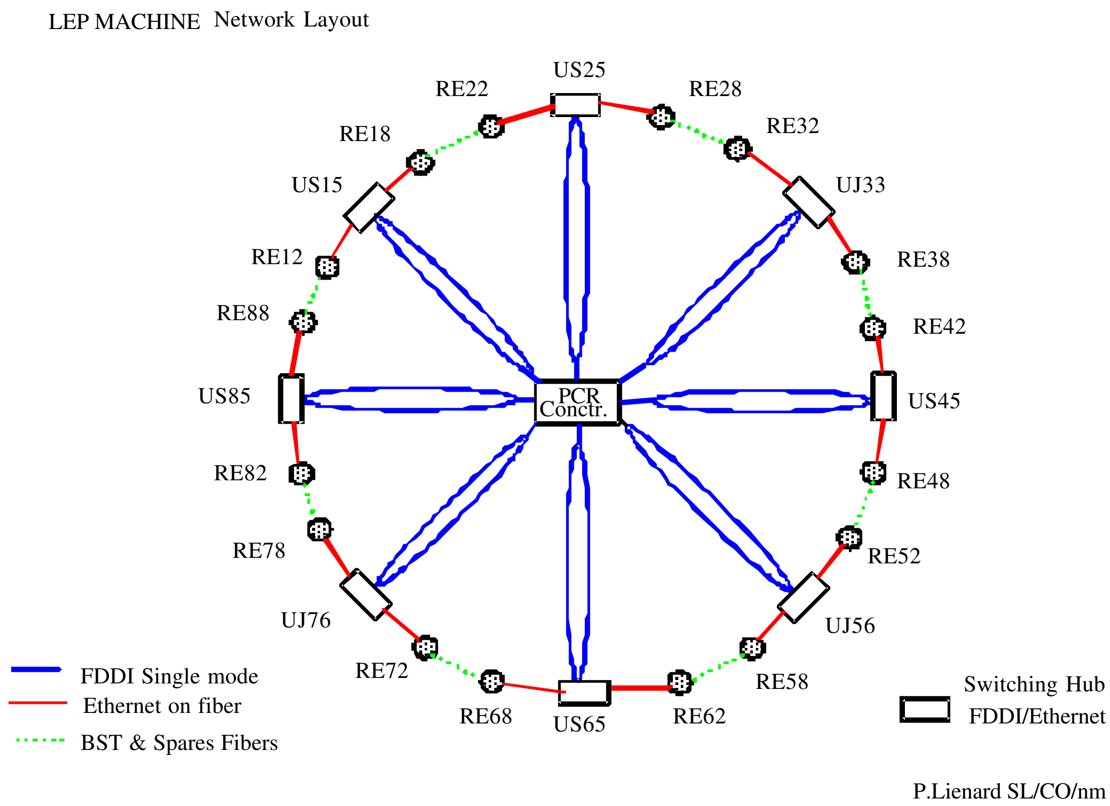

Figure 1: New LEP Machine Network

Current planning is to extend this work to the other LEP backbone, used for the control of the LEP Services. The installation of the associated fibres and electronics is an investment in the infrastructure that will be used for the construction and commissioning of the LHC. After LEP the SPS will also be upgraded to a 100 Mbit/s backbone. Prototype work has already started in the SPS [6] for a new generation of multi-turn orbit acquisition electronics. This equipment will be used during study sessions aimed at preparing the SPS as an injector to the LHC. Such developments place ever increasing demands on the available networking bandwidth.

Operational experience with the SL network has shown us the very stringent requirement for high availability and reliability. Network management tools are essential and we are using the HP OpenView [7] product. Also essential for the reliable operation of the SL Control System is the ability to reboot any remote computer, bridge router or gateway via a separate network. This outband network is based on long distance MIL-1553-B fieldbus links working at 125 kbit/s.

3 FRONT-ENDS

The LEP front-end machines have provided the required functionality to fulfil their role as generic equipment access machines. There have been reliability problems, initially caused by disk failures; disks were replaced but, in recent years, fault rates are still high but distributed across all components. In addition these machines are running XENIX which we cannot upgrade and which is increasingly difficult to support as there are no NFS, TCP/IP sockets, or SNMP agent. Over 150 VMEbus front-ends have been added to LEP to enhance the performance of specific systems. The majority of Beam Instrumentation and RF front-ends are implemented in VME and OS9.

After a common study with the PS Division [8], it was decided in 1991 to base further evolution on diskless LynxOS systems, based either on PCs or VMEbus crates. The SPS [9] is equipped with about 50 diskless LynxOS 486 PCs as the general purpose front-ends; Figure 2 shows a typical installation. A vertical bus extension is used to connect a VME chassis housing the bus controllers for the MIL-1553-B field bus. Such an architecture follows the LEP model but we will simplify these machines at the beginning of 1996. Final tests are being carried out on a new PC bus MIL-1553-B bus controller which will eliminate the intermediate VME crates and associated modules and software. As with LEP, high performance specific front-ends for the SPS are VME based but many of these systems run LynxOS. There are about 30 such systems.

Following the study with the PS Division we are gradually replacing a central file service with a set of regional servers to act as boot servers and fileservers for the front-end machines. These servers will emphasise the isolation along the same frontiers as the sub-networking. The SPS Machine Network is equipped with an HP-735 server for this purpose.

Last year a project [10] was launched to upgrade the LEP XENIX PCs. The hardware solution selected makes use of the VME chassis associated with each of these machines. The PC will be discarded and replaced by a processor card in VME with an on-board Ethernet interface. This choice has the advantages of simplifying the front-end architecture, of approaching the configuration used by many equipment groups in specific front-end machines and of moving to more robust hardware. Naturally, the new machines will run the LynxOS operating system. A market survey for information on VME-based processor cards running LynxOS indicated that the PowerPC processor family offers the best chance of long term commercial support. These processors have a high performance and are of interest for other applications within the SL Control System. Although the primary benefit from these new systems will be improved reliability and eased maintenance, the PowerPC 603 chip will also offer considerably higher performance. Its rating is 60 SPECint92 and 70 SPECfp92 with a 66 MHz clock, and the new PowerPC 603e, which is 603 pin compatible, offers 120 SPECint92 and 105 SPECfp92 with a 100 MHz clock. First deliveries to CERN have just been made.

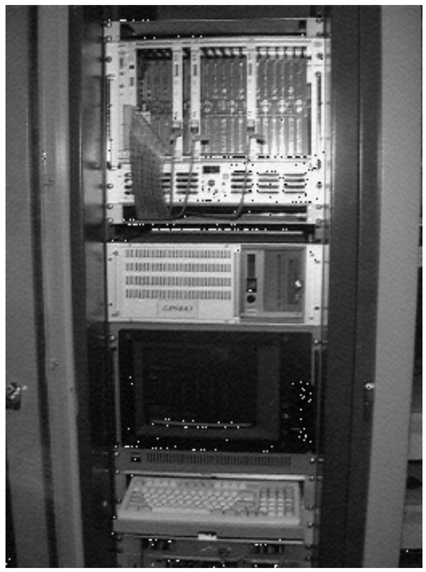

Figure 2: SPS Front-End

Figure 3: PCR Control Desks

After the shut-down of LEP around the end of the decade, we expect this new RISC processor family to remain in service. It will continue to ensure the control of the primary service equipment around the LEP/LHC ring and be used to upgrade the performance of SPS equipment. As we have decided to use the VMEbus standard for the front-end machines it will be easy to enhance their computing power to match the future LHC requirements. The 603 processor can be exchanged for the 604 or the VME processor card can be replaced. Another option will be to connect the front-ends directly to the FDDI network via a PMC-FDDI interface.

4 BACK-ENDS

The use of commercially available single screen workstations for LEP control in 1989 was a fairly novel choice. Of course with the power and accessibility of modern windowing software the operational experience has been entirely positive. All of the control software for LEP was implemented on Apollo workstations making full use of the excellent graphics and networking systems. From 1991, we adopted a policy to change the operators' workstations to 19" colour X Terminals hosted by HP-700 series HP-UX servers. Since that time, a large majority of new application software has been written for this environment. The application software for SPS and LEP is sufficiently stable that we will have to carry out a specific migration project to eliminate the obsolete Apollo workstations in the coming years.

During the SPS upgrade we installed X Terminals in the remote equipment building of the SPS, and the possibility to run "Control Room" software next to the equipment has been appreciated by the teams responsible for commissioning and trouble shooting. These X Terminals are hosted by the regional server mentioned in the previous section. Our aim is to provide the same environment for application software in the Control Rooms, remote buildings and at home. We have installed two layers of workstation screens and X Terminals at the control desks of the SL operators (see Figure 3). The upper layer is dedicated to "Fixed Displays" presenting performance, diagnostic and alarm data.

The Fixed Displays are generated either on X Terminals or via a teletext generation system. Teletext is the preferred solution if a display is to be made widely available on the site. An alternative access to the teletext displays is provided via an interface to an http server, presently only accessible to machines at CERN. All the information for these displays is collected from the SPS and LEP by a pair of dedicated central servers. The data are also saved in Oracle and used for performance statistics and machine physics studies [11]. We run Oracle on a pair of Sun SparcServer 1000 machines in the CERN computing centre.

As modern software becomes more and more complex we are continually upgrading the resources to support these products. As we turn to commercial tools to reduce our own application development effort the demands increase very quickly. We are currently upgrading to a twin CPU HP-UX K Server to host the 15 X Terminals in the PCR.

For the future we expect to see a continuation of the trend towards fewer, more powerful, central server machines. The resources will not be defined by the accelerator modelling activities in the Control Room but by the ever increasing CPU and network bandwidth for communication requirements. We are also expecting a role of growing importance for fully mobile terminals. Tests have started with both GSM telephone- and radio-based networks.

5 APPLICATION SOFTWARE

Some experience of software engineering practices was gained during the production of the first generation of control software for the LEP Control Room. With the introduction of a commercial workstation as the machine console it was possible to use commercial products both for the development and implementation of software. SASD tools were used in the design of some of the machine applications and an industrial drawing package was used for the synoptic views and the control of all the LEP Services equipment.

The trend over the past years has been towards greater use of these tools. A second generation of control software for LEP [12] relied heavily on data modelling, database design and a commercial RDBMS, Oracle, to rationalise the data management involved in operating the accelerator.

Pushing the use of industrial software engineering standards further we have carried out a full-scale exercise [13] to sub-contract the production of a Control Room application program to an external company. As this was the first experience with such an approach at CERN there was an extended project definition phase during which we decided to draw upon the experience of the aerospace industry and base the project management on the ESA PSS-05 software engineering standards [14]. A competitive tender was prepared using these standards asking for the development and installation of software following a specification based on the users' functional requirements. The successful bidder undertook all the development on his own premises apart from a series of project review meetings and the final installation. Several important conclusions were derived from this exercise:

- It is possible to write a contractual specification for well-understood applications,
- The development can take place remotely,
- A considerable effort is required to maintain control of the project,
- Specific outsourcing overheads are small, given that a formal development is used.

Perhaps the most unexpected difficulty with this project came from organising the installation and testing schedule to adapt to the availability of the accelerator. We will continue to develop our experience in this domain.

The emergence of X Windows and OSF/Motif as graphical standards has encouraged us to define a rich development environment, the X Windows User Interface Management System [15] for the generation of operational graphical user interfaces required to control the SPS and LEP and the CERN wide technical services. These graphical user interfaces are running in both PCR (SPS and LEP Control Room) and TCR (Technical Services Control Room) and are also distributed in several other Control Rooms (SPS North Area Control Room, PS Control Room). XUIMS is a modular development environment allowing the interactive creation of OSF/Motif graphical user interfaces. Fully compatible with the OSF/Motif standard, it is regularly enriched with new widgets developed by industry. XUIMS is based on the OSF/Motif widgets set and the X Toolkit Intrinsics (Xt) which is the standard mechanism on which many of the toolkits written for the X Window System are based. This environment is intended to ease the development of graphical user interfaces including interactions with 2D and 3D scientific data and access to SPS and LEP equipment. Several libraries and widgets are available to test applications interactively and to reduce the C coding effort to a minimum.

A recent development in the software industry has been the availability of “middleware” packages encompassing the software infrastructure requirements to bridge the gap between operating systems and application software. Such a package offers obvious attractions to us as this middle ground has consumed manpower resources in the past and it has been difficult to achieve a homogeneous software environment for applications produced at different stages of the evolution of our system. A full-scale application of one of these packages is being prototyped [16] by the Technical Services team at CERN and we are evaluating the use of the “middleware” product as a standard application software environment for the future.

Despite these attempts to harness commercial products and to outsource software development we believe that operation of our accelerators could still benefit from better application software. Our difficulties to achieve an optimum management of this activity are described elsewhere at this conference [17].

6 INTEGRATION OF INDUSTRIAL SYSTEMS

During LEP construction certain equipment groups chose to satisfy their control requirements with turn-key industrial Control Systems. These groups were responsible for the technical services of the LEP machine, and by that time there was a variety of manufacturers offering such products. As the philosophy of the Control Group was to hand over responsibility to the Equipment Group at the ECA layer [1], these systems were interfaced by links, often simple serial lines below the ECA. This arrangement has the advantage that, seen from above, these systems are connected by the standard MIL-1553-B fieldbus, and the software in the ECA presents the user with a view of the system which follows the general architecture of the Control System. The disadvantage has been that the supervisory software available from the manufacturer cannot be used from the central Control Rooms. Generally, groups who adopted this solution made only small investments in Control System expertise within the groups. This has led to persistent difficulties for these groups to provide adequate local control facilities and to maintain their software.

To rejuvenate old and obsolete technical systems, industrial equipment is increasingly being installed and new projects for LHC are now taking shape. In order to smoothly integrate new industrial systems into the Accelerators and Services Control System, strategies are evolving [18] for interfacing this equipment, for supervising it centrally and for operating it locally during commissioning and maintenance.

As the architecture of the Control System is modelled on three layers, Control Room, Front-End Computing and Equipment Control, industrial equipment and systems can be interfaced at each of these levels.

The method of interfacing and integrating this industrial equipment depends on several parameters, such as its type (PLC or workstation), its openness to International Standards, the size of the process to interface, the need for local and/or global communication, the required process real-time response, and so on.

Interfacing at the Control Room Level

In the case of the delivery of a complete industrial turn-key system, all the hardware and the software are provided by the contractor according to a CERN Functional Specification. The manufacturer’s supervisory station is located in the Central Control Room and CERN operators have to learn how to operate this proprietary system and become acquainted with its particular man-machine interaction and behaviour. Private connections from any local equipment to the central Control Room have to be established via dedicated cables, fieldbuses, TDM channels or private networks; all equipment supplied is to be installed, commissioned and maintained by the contractor.

A “gateway” between the manufacturer’s supervisory station and the CERN control network has to be provided and programmed in such a way as to allow a limited high level remote access and for integration into central services such as the alarm server and the logging database. The system must offer remote control from X Terminals on the SL Control network. The SPS was equipped with a new access Control System at the beginning of 1995 following this integration scenario. The advantage of such a complete industrial turn-key system is that the manufacturer provides all the hardware and writes the software; he has the full responsibility for the project and brings to CERN his “expertise” in his domain of activity. The drawback is the duplication of the whole control network; this duplication has to be repeated with the installation of each new industrial system of this type. Every manufacturer provides the operators with a specific man-machine interface, and the supervision of the system is not integrated with other systems.

Interfacing at the Network Level

Complete industrial turn-key systems can be connected directly to the control network provided they have an Ethernet interface, use TCP/IP for remote communication, work in the UNIX environment, support MOTIF and X11 as their graphical interface and allow SNMP to supervise their network connection. An increasing number of manufacturers offer, or will offer in the near future, this openness to International Standards.

If the manufacturer's protocol conforms to an International Standard or can cohabit with other private protocols and TCP/IP on the CERN control network, then the existing network infrastructure can be used avoiding expensive duplication. The LEP and SPS networks are being upgraded [5] to FDDI, and the associated electronics uses switching routers so that a multi-protocol networking service will be provided. As central and local control rooms are equipped with X Terminals, industrial equipment can in this case be accessed widely using the supervisory software provided by the manufacturer.

Interfacing at the Front-End Level

Industrial equipment has been integrated by connecting industrial controllers to the CERN LynxOS front-end computers using the manufacturer's preferred fieldbus. At present, LynxOS drivers are available for GPIB/IEEE-488, BITBUS, RS-232/422 and MIL-1553.

Additional drivers for other popular fieldbuses (PROFIBUS, SINEC-L2, WorldFIP, FIPIO, mini-MAP, etc.) can be made available for new industrial applications. The advantage of interfacing at the front-end level is that the Control System can accept a wide variety of industrial equipment fully integrated with CERN-made systems and into the software environment such that the full functionality of the equipment is available to the high level accelerator control software. The disadvantage is that the turn-key solution is no longer possible. CERN must be heavily involved in the production of the drivers and local and central supervisory software.

Interfacing at the Equipment Level

This is the lowest level where industrial equipment can be interfaced, either by direct connection to a PC or a VMEbus crate or else via one of the SL/CO supported fieldbuses. Most industrial PLCs provide at least one RS-232/422 connection in addition to their private network. Interfacing via RS-232/422 links is supported both on the front-end directly and on the MIL-1553 fieldbus via an RS-232/MIL-1553 interface box. The connection to the fieldbus can only be justified for isolated equipment geographically dispersed over the CERN site or inside an accelerator building.

7 CONCLUSIONS

The LHC Control System will be a combination of specialised high performance systems for the monitoring and control of the intense proton beams and industrial Control Systems responsible for the large amount of service equipment distributed around the 27 km ring. All these systems must collaborate to achieve the high performance and reliability that will be required of the future machine. Such an orchestration will be achieved by appropriate integration of these systems into the control and monitoring facilities available to the accelerator engineers and physicists in the future LHC Control Room. Our concern is to hit the right balance between autonomous turn-key solutions which take full advantage of the services offered by industry and more specific developments offering very flexible functionality but necessitating heavy involvement by CERN in the technical development and maintenance.

8 ACKNOWLEDGEMENTS

The evolution of the SL Control System is the collective effort of the SL Control Group and control specialists in various groups in the SL and other divisions at CERN. During much of the period covered in the paper the SL Control Group was lead by K.H. Kissler and it is a pleasure to acknowledge his contribution.

9 REFERENCES

[1] P. G. Innocenti, "The LEP Control System", ICALEPCS'89, Vancouver, BC, Canada, October 30-November 3, 1989. Proceedings Nucl. Instr. and Meth. in Phys. Res. A 293 (1990) 1-5.

[2] K.H. Kissler and R. Rausch. "New Control Architecture for the SPS Accelerator at CERN". ICALEPCS'91, KEK, Tsukuba, Japan, November 11-15, 1991. KEK Proceedings 92-15 (Dec. 1992) 53-58.

[3] P. Liénard, "The SPS and LEP Control Network Architecture", ICALEPCS'89, Vancouver, BC, Canada, October 30-November 3, 1989. Proceedings Nucl. Instr. and Meth. in Phys. Res. A 293 (1990) 215-220.

[4] P. Liénard, “Evolution of the SPS and LEP Communication and Control Network for the LEP and SPS Accelerators”, ICALEPCS’93, Berlin, Germany, October 18-23, 1993. Proceedings Nucl. Instr. and Meth. in Phys. Res. A 352 (1994) 170-172.

[5] P. Liénard, L. Guerrero, “A new LAN Concept for LEP Machine Networks, a Step towards LHC”. This Conference, ICALEPCS’95, Chicago, ILL, USA, October 30-November 3, 1995.

[6] Private Communication. H. Schmickler, CERN.

[7] HP OpenView, Network Node Manager, Hewlett-Packard Company.

[8] R. Rausch and Ch. Serre, (eds), “Common Control System for the CERN Accelerators”, ICALEPCS’91, KEK, Tsukuba, Japan, November 11-15, 1991. KEK Proceedings 92-15 (Dec. 1992) 53-58.

[9] P. Charrue and M. Clayton, “The New Controls Infrastructure for the SPS”. This Conference, ICALEPCS’95, Chicago, ILL, USA, October 30-November 3, 1995.

[10] A. Bland, P. Charrue, F. Ghinet, P.Ribeiro, M. Vanden Eynden, “The new Generation of PowerPC VMEbus Front-End Computers for the CERN SPS and LEP Accelerators Control System”. This Conference, ICALEPCS’95, Chicago, ILL, USA, October 30-November 3, 1995.

[11] R. Billen, E. Hatziangeli, I. Laugier, I. Reguerro, N. Segura, “Mission Critical Database for SPS Accelerator Measurements”. This Conference, ICALEPCS’95, Chicago, ILL, USA, October 30-November 3, 1995.

[12] R. Bailey, A. Belk, P. Collier, M. Lamont, G. de Rijk, M. Tarrant, “Development of the LEP High Level Control System using Oracle as an On-line Database” ICALEPCS’93, Berlin, Germany, October 18-23, 1993. Proceedings Nucl. Instr. and Meth. in Phys. Res. A 352 (1994) 430-433.

[13] B. Denis, “Outsourcing the Development of Specific Application Software using the ESA Software Engineering Standards: The SPS Software Interlock System”. This conference, ICALEPCS’95, Chicago, ILL, USA, October 30-November 3, 1995.

[14] ESA Software Engineering Standards, ESA PSS-05-0 Issue 2, Board for Software Engineering Standardisation and Control, ESTEC publications (1991).

[15] M. Vanden Eynden, “The X-Window User Interface Management System (XUIMS) at CERN”. This conference, ICALEPCS’95, Chicago, ILL, USA, October 30-November 3, 1995.

[16] P. Ninin, H. Laeger, S. Lechner, R. Martini, D. Sarjantson, P. Sollander, A. Swift, “The Technical Data Server for the Control of 100’000 Points of the Technical Infrastructure of CERN”. This conference, ICALEPCS’95, Chicago, ILL, USA, October 30-November 3, 1995.

[17] R. Bailey, “How can Operations get the Application Software they want?”. This conference, ICALEPCS’95, Chicago, ILL, USA, October 30-November 3, 1995.

[18] P. Anderssen, P. Charrue, R. Lauckner, P. Liénard, R. Rausch, M. Tyrrell, M. Vanden Eynden, “Interfacing Industrial Equipment to CERN’s Accelerators and Services Control System”. CERN-SL-95-42 CO Report, April 1995, Geneva, Switzerland.