Overview of Beamline Control and Data Acquisition at ESRF

F.Epaud on behalf of the ESRF Computing Services .
ESRF
BP 220, F-38043 Grenoble Cedex
FRANCE

Abstract

The ESRF provided photons to beamlines at the end of 1992. Two and a half years later, 14 beamlines (+ 4 CRG beamlines) are open to users and regularly receive visiting scientists. 6 new beamlines are in their final commissioning phase and will be operational early in 1996. This rapid building schedule has been achieved by different small teams using exactly the same technologies for all the beamlines, including the CRG (Collaborating Research Groups). VME front-end hardware and software originally designed for the machine control system, based on the so called ‘device server model’, have been reused. This model uses client/server communication via RPC and runs on OS-9 (and LynxOS) on the VME front-ends and Unix workstations. In the data acquisition domain, we are currently extending this model to define a ‘Modular Data Acquisition Software’ to handle the acquisition memories (AM) that are used to collect events or images. The main parts of this system are the ‘online display’ used to visualise and evaluate rapidly the quality of the raw data collected inside the AM, and the ‘fast data transfer’ which sends large amounts of data (of the order of several hundreds of mega bytes) at high speed to a central computer facility (NICE) which has a large storage, computing and soon graphics capabilities. The transfer can be effected over Ethernet or ATM. The centralised computing facility (NICE) has been set up to allow short-term data archiving and data treatment by means of dedicated file servers using RAID disk storage arrays, optical and magnetic tape robots for data migration and backup and a loosely-coupled workstation cluster for data analysis and number crunching.

1 Introduction

It is planned to install 30 ESRF beamlines open to the public by the end of 1998 to which 12 Collaborating Research Group (CRG) beamlines will be added [1]. Today, 14 ESRF and 4 CRG beamlines are operational and regularly receive visiting scientists. By the end of this year, 6 new ESRF beamlines will be commissioned, which will increase the total number of open beamlines to 20 and the number of end stations to 24.

The choice made by the council to commit the ESRF to build the beamlines instead of having too many built by the CRG, as it is the case at many other institutes, allowed us to use the same technology on all the beamlines and enables rapid provision of usable instruments to the scientific community. Most of the CRG beamlines also use the ESRF technology.

Both the beamline control and data acquisition use the so called 'Device Server Model' derived from the 'Standard Model 91' defined at ICALEPCS 91, and originally designed for the 'Machine Control System' [2, 3]. This model has been extended for data acquisition needs by the definition of a 'Modular Data Acquisition Software' [4, 5] to handle the acquisition in a standard way for most of the ESRF systems.

On the back-end, the ESRF also provides a centralised computing facility (NICE) for short term data storage and data analysis. Transfer of data is normally made over Ethernet but ATM is also used for beamlines providing large amounts of data.

2 Beamlines architecture

Each ESRF beamline has its own control system and its own network, in order to run autonomously. A router allows communication with the rest of the ESRF and the outside world and also allows access to the 'Machine Control System' through a gateway, to get parameters like machine current intensity, insertion device setup, etc. A dedicated X terminal is connected to the private machine network to control the insertion devices.

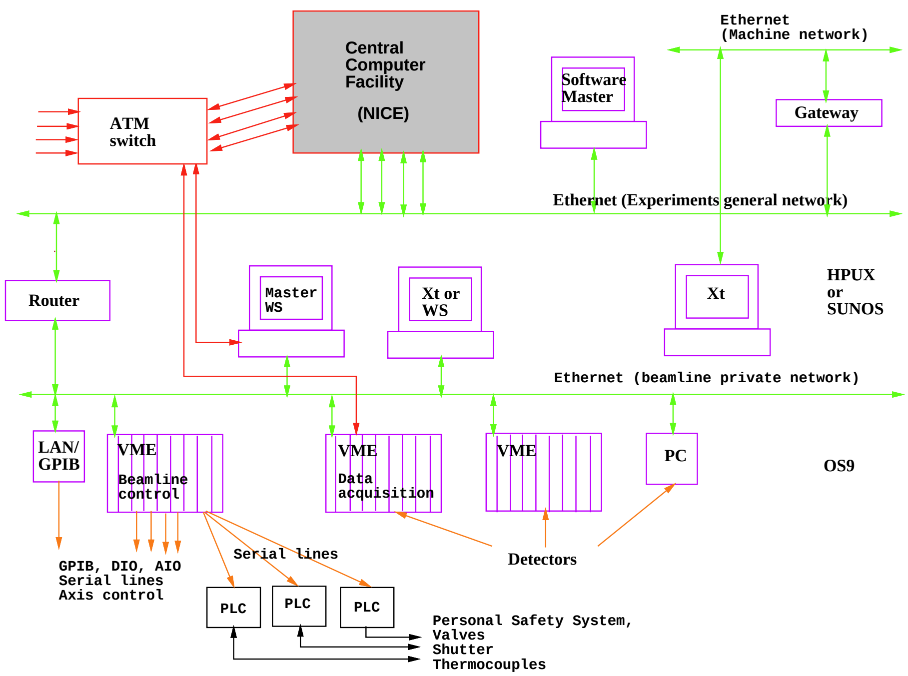

Figure 1: Beamline architecture

A typical beamline is composed of a main UNIX workstation (HP700, Sun) running X11/Motif GUI or a CLUI (Command Line User Interface) client like SPEC [8]. The workstation is used for the beamline control system and often for data acquisition. Other UNIX workstations or X terminals are also added for preliminary data analysis.

On the lower level are the VME systems using the Motorola MVME167 68040 CPU running OS-9, used for instrument control.

Generally at least one VME crate is used for the control of the optics components, vacuum, slits, attenuators, etc. and at least one other crate is used for data acquisition. For beamline control, these crates drive a large number of serial lines, digital or analog inputs/outputs and axis control.

We also use many GPIB devices by means of LAN/GPIB converters directly connected to Ethernet. This choice has proven to be the most reliable one.

The security aspects are left to the PLCs, which are in charge of the vacuum and Personal Safety System interlocks. They are accessed by the VME via serial lines.

PCs are also used as stand-alone systems, mainly to run commercial acquisition systems, such as multichannel analyzers, CCD cameras, image plate scanners, etc.

The philosophy adopted at ESRF is not to keep the collected data on the beamline system but to transfer them to a central computer facility (NICE: Network Interactive Computer Environment). This can be reached by the experiment Ethernet backbone network through the router. However, for beamlines collecting large amounts of data, the transfer rate was too low and overloaded the network. Therefore ATM optical links rated up to 155 Mbits/s have been set up. On some beamlines which use large histograming memories in VME, the crate is directly connected to ATM by means of a VME/ATM adaptor.

Last but not least, a software master to compare beamline software against a reference has been set up. This recently installed mechanism is of great help to keep all beamlines at the same software level and to detect potential problems at an early stage.

3 The Modular Data Acquisition Software

The central part in most of the ESRF acquisition systems is the acquisition memory (AM) into which raw data are acquired. Usually this memory is in the front-end (VME or VXI crate) and has a typical size of the order of several hundreds of megabytes.

Processes dealing with this memory can be divided into two groups. In the first group are so-called producer processes, which define memory organisation and put data into the AM. In the second group are so-called consumer processes, which nondestructively access the AM.

The data structures reflecting the organization of the AM are kept in another memory area usually located on the CPU card, and shared among the processes involved. It is called the AMS (Acquisition Memory Structure).

To minimize the amount of software to be written for different acquisition systems we have defined an API [6] (MEMAPI, MEMory Application Programming Interface) which allows the software to access memory structures (partitions, images, ROIs, pixels) in a uniform way, in total abstraction of the physical arrangement of the data within these memories. Another part of the software dealing with the AM is the data access layer [7] (DATA-ACC) which is the only hardware dependent part needing to be rewritten or ported to every new acquisition system.

Data acquisition control processes that are data producers are processes that are specific for each acquisition system, and in most cases they are written in the form of a device server. Client processes which communicate with the acquisition control process usually run on a UNIX workstation. Their interactive user interface can be either in the form of an alphanumeric menu, X11/Motif graphics users' interface or SPEC application [8]. Additionally we have written two general consumer applications, which are the online display, and the data transfer process.

In order to visualize data while they are being acquired (histogrammed, listed or accumulated), the online display process runs in the front-end CPU. This is an X11/Motif client [9] which, in its current version, allows the display of 2D image data in snapshot or live mode, can show successively a series of images in a cyclic fashion, can pan and zoom, do Z-slicing and X or Y cuts (1 or more pixels wide). Typically, we achieve in live mode a rate of 1 image/second using an MVME167 CPU running OS-9 and an X-server on an HP or SUN workstation. We are also currently evaluating a solution based on an ELTEC-EUROCOM17 bi-processor card (2 x 68040) having embedded graphic capabilities which allow the X11 server to run locally and access the Video RAM via external hardware to refresh the image in real time.

During or after the acquisition run, data are transferred to a central computing facility. The data transfer process, written in the form of a device server, uses the XFRLIB library [10], which contains functions manipulating BSD sockets above TCP/IP. On the physical level, we use Ethernet as well as ATM, either on workstations or on VME crates. For the VME connection to ATM we use a dedicated MVME167 CPU running LynxOS and handling TCP/IP and the ATM driver.

Due to the system modularity and the tightly coupled communication mechanism based on VME shared memory, the processing power can be distributed over several processors. For example, the acquisition system control could be performed by a MVME162 CPU located in slot 1 (which is also in charge of the crate control), a second ELTEC E17 CPU with OS-9 could be used to run

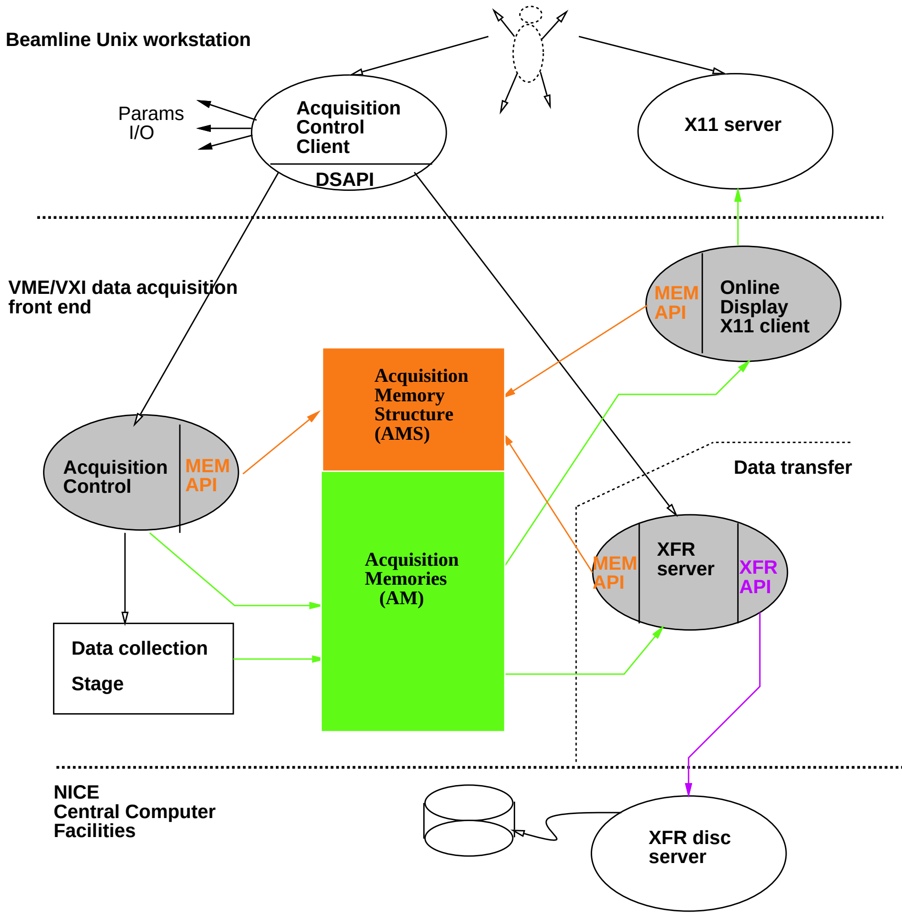

Figure 2: Software architecture of a data acquisition system.

the the X11/Motif online display client and a third MVMEM167 CPU could be used for the data transfer running LynxOS, the TCP/IP network stack and the ATM driver.

This architecture is now implemented or under implementation on all the data acquisition systems designed within the ESRF Computing Services, such as the X-ray image plate scanner, Multi-Wire Chamber Gas-Filled Detector [5], CCD camera, etc. Of course, ESRF beamlines also use commercial systems controlled by PCs or UNIX workstations.

4 NICE: Networked Interactive Computing Environment

The large amount of data collected by the beamline end stations are stored after being transferred, and can be analysed on the central computer facility (NICE) [11].

Due to the volume of data, it is not possible to archive all of them, therefore we have been obliged to fix a limit of 100 days, after which the data are automatically deleted. In addition, visitors normally go back to their home institutes with their data and delete them on NICE before leaving the ESRF.

The 100 days policy is achieved by a three level file migration facility providing a total of 1.5 tera-bytes storage capability consisting of:

- 200 giga-byte RAID7 disk arrays.
- 180 giga-byte of rewritable optical disk.
- 1200 giga-bytes of Exabytes tape robot.

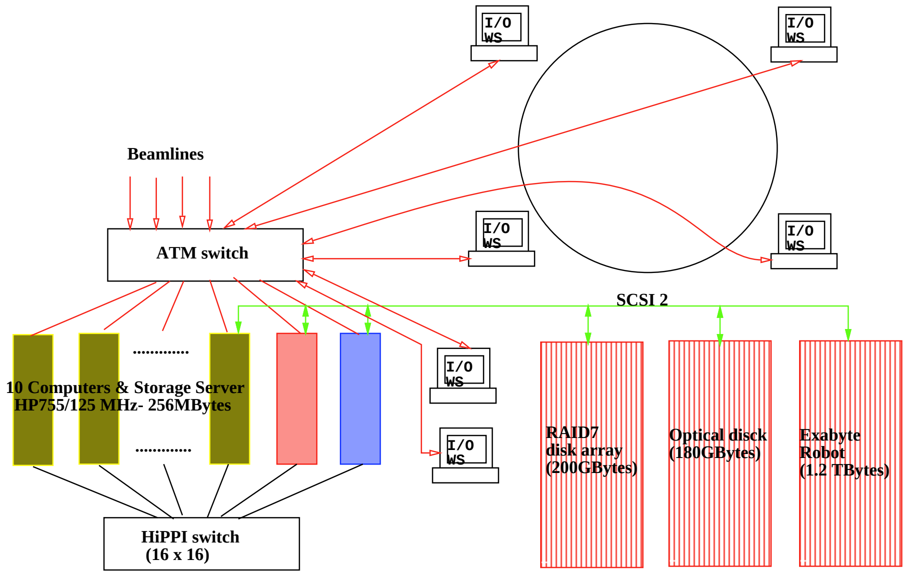

Figure 3: NICE architecture

Data analysis can be performed on a HiPPI cluster of 10 HP755-125, each equipped with 256 mega-bytes of RAM which can be used individually or as parallel computers.

All these computers, including the ones used to control the storage system, are linked with a HiPPI bus (High Performance Parallel Interface) with a bandwidth of up to 800 Mbits/s. Each computer is configured such that a user sees exactly the same environment when he logs on any one of them. Powerful graphics capability and an increase in the number of stations on the cluster are currently under evaluation.

Six I/O workstations, connected with a 155 Mbits/s ATM links to NICE are also available around the ring to allow the users to make their own DAT or Exabyte tapes before leaving the ESRF.

5 Conclusion

The rapid building of the 14 first beamlines has been made possible by reusing the improved hardware and software technology developed for the 'Machine Control System'. We are currently working hard to install the 16 new beamlines scheduled to be operational by the end of 1998, which will all use the current architecture. Even though a great deal of manpower is dedicated to beamline installation, we are constantly improving our 'Object in C' model and are currently introducing C++. Investigations are under way to use ORBIX from Iona (based on CORBA from the OMT) and/or OLE from Microsoft in order to migrate to a system of distributed objects.

6 Acknowledgements

The author wishes to acknowledge the contributions of all the members of the Computing Services.

References

[1] ESRF Highlights 1994/1995. ESRF internal report, September 95.

[2] A.Götz, W.D.Klotz, J.Meyer: Object Oriented Programming Techniques Applied to Device Access and Control in International Conference on Accelerator and large experiment Physics Control System, Tsukuba, November 1991.

[3] A.Götz et. al: Experience with a STANDARD MODEL' 91 based Control System at the ESRF. Proceedings of the International Conference on Accelerator and Large Physics Experiments Control Systems. Berlin, November 93.

[4] F.Epaud, F.Sever, A.Beteva: Modular Data Acquisition Software. ESRF internal report, reference: DASR006, January 1995.

[5] F.Sever, F.Epaud, F.Poncet, M.Grave and V.Rey-Bakaikoa: Modular Data Acquisition System and its Use in Gas-filled Detector Readout at ESRF. Synchrotron Radiation Instrumentation 1995 (SRI95), October 95.

[6] F.Sever, F.Epaud and M.Grave: Memory Library Application Programmers' Interface (MEMAPI). ESRF Internal documentation, reference: DASS003. 1995.

[7] M.Grave: Data Access Programmers Interface (DATA-ACC). ESRF Internal documentation, reference: DASS007, 1995.

[8] Certified Scientific Software: Xray Diffraction Software. Cambridge. MA.

[9] F.Poncet: Online Display Users Guide. ESRF internal report, 1995.

[10] F.Epaud: Sockets Based Data Link Users Guide (XFR). ESRF Internal report DASUG001, 1994.

[11] R.Dimper, B.Lebayle, WD.Klotz: Technical Specification for a Central Computing Facility. ESRF Internal report, October 1993.