

Computing, Control, and Data treatment at the ILL 1974-2008

When the British joined the ILL in 1973 a new department of Computing and Electronics, DIE, was created with a British head of department, Dennis Rimmer, with the responsibility to develop the control of instruments and the treatment of the subsequent data. The secretaries included Christianne Tillosen (replaced at times by Winnie Hollenburg). Initially Axmann was in charge of electronics and control computing, with Taeschner in charge of the computers. Axmann returned to Germany in 1975, shortly followed by Taeschner. Drexel was given charge of the electronics group responsible for the counting and positioning systems, and the control computers. The central computer group, Calcul Scientifique, was headed by Siret, which operated the PDP10 central computer, and provided general support including numerical analysis. During the long shutdown an FNE plan was implemented and Rimmer was forced into retirement. The Institut was reorganised with each deputy director becoming head of a division. Computing and electronics became groups within the Projects and Techniques division. Scientific computing activities became a group in the Science Division; fortunately, apart from formal management, the cross-collaborations between scientists and the technical groups remained active.

At the initial stages of the Institut there were few operational instruments, and the central computer was not available. The original choice of control systems had been made a couple of years earlier, and were strongly based on using shared minicomputers, although already experience in the UK at Harwell, and in the US, showed the complexity of such systems and unreliability were unfavourable compared to the increased cost of individual computers. The committee of two experts from France and Germany chose a French and a German control computer manufacturers to control the step-scanning diffractometers and three-axis spectrometers, and multichannel data acquisition for a second set of instruments, respectively. When Mössbauer replaced Maier-Leibnitz as director he knew of the utility of the PDP10 computer in the nuclear physics community, (and also that it could offer real-time facilities to experiments). By proposing the rental of such a system he obviated the nationality requirements for a major acquisition. It was also possible to install this machine in not much more than a corridor!

Summary of Systems Evolution

This is a subject which has always been open to management decisions based on political sensitivities rather than practicalities. The control computers were chosen by a committee of one German and one Frenchman; each computer served up to 6 instruments. Until Mössbauer there was no Central Computer (1973). Once installed there was reluctance to enhance the computing power though crystallographers at the Grenoble CNRS group appreciated the time-sharing facilities it offered. Users were encouraged to take away data even where the instrumentation and data treatment packages were specific to the ILL. A few years

later at Saclay the instruments were controlled by individual Honeywell Mini-6 computers. Elsewhere individual DEC PDP8 and PDP11 computers were the norm eliminating many conflicts evident on shared systems. By 1983 at RAL individual DEC VAX11/730 32-bit computers were used from the start on the 8 or so instruments, all with the same acquisition electronics first on the computers UNIBUS; later this was attached through an SCSI interface on the microVAX. The greater number and diversity of instruments at the ILL always precluded a single solution. Already by 1978 the inefficiency of existing ILL systems was clearly evident (the D8 diffractometer typically only counted for 25-40% of time, awaiting controlled positioning) and their replacement was being planned.

To aid the scientific program a small group, GRATIN was created (Filhol, Hueskin, Ghosh) with knowledge of activities in Diffraction, Nuclear Physics, and condensed matter physics respectively.

As part of the Deuxieme Souffle (1978) the CARINE and NICOLE systems would be replaced, and the Central Computer was to be upgraded. For instrument control there was strong support for the European industry, and seven SEMS-Solar 16-40 computers were acquired for D1B, IN1, IN2, IN3 etc. While the hardware outperformed the PDP11/34 systems chosen for the NICOLE instruments, IN4, IN4, D7, and PN1, the poor software support and specialised electronics interfacing required resulted in years of delay for the Solars. Only four of the seven systems acquired were ever put into service. There was little feedback from the scientists since at that time there were few data treatment programs running on minicomputers. The task of creating control functions required programming in close collaboration with the electronics engineers.

Central Computer Replacement 1979-91

A more heavily political choice was that to replace the DEC-PDP KI-1070, by 1980/81. Already over seven years old it was heavily used by all groups. This project is described in more detail in Appendix I. The new computer, a DEC KL1091S was finally installed in a new computer building in 1981, and was essentially a faster version of the original KI-1070, with better networking front end processors, and new tapes (6250bpi), Printers (1200lpm) etc. Some of the KI10 peripherals were retained and reused. Optimem, 12 inch re-writeable optical disks were used for data archiving up to 3Gb of data, essentially tape image files, reducing pressure on the 9-track tape drives.

As the PDP11, with its general purpose UNIBUS, evolved into a single board computer (LSI11/PDP11/23 using a cheaper multiplexed Q-bus) it became much cheaper and it was offered for sale OEM by Plessey Ltd, a British company. It could be obtained without too much political bargaining and progressively replaced instrument computers, mostly running the RSX11M system, though RT11 and TSX (a multi-users version of RT11 were also used). By the mid 80's 32-bit computers were dropping in price. A VAX11/750 was bought to manage data for the multi-detector single crystal diffractometer D19, which was controlled by a PDP11/24 (a UNIBUS variant of the LSI/11). VAX11/730 computers were installed to control D11 (1985) and D17 (1986) directly without a front-end PDP11. The advent of the microVAX-II single board computer with a Q22-bus lead to price falls and it became very attractive to replace the PDP11/34 bought in 1980, by these more powerful systems.

Apart from deliberately slowing response times in the software for faster cpus, the low level VMS software written for D11 was used unchanged on all subsequent VMS instrument computers.

The network was developed to include ethernet to all instruments. The introduction of a VAX8600 in 1986 to take over from the DEC10, then offered complete compatibility with 32-bit instrument computers, again using the ethernet local area network with DECnet throughout. The VAX was popular with scientists, and to aid use by the ESRF the 8600 was upgraded with a financial contribution from the ESRF, to a VAX8650. The actual configuration was a cluster including an HSC50 storage controller, and a VAX11/751 (recovered from the upgraded PDP11/55 NICOLE replacement concentrator). It was hence easy to expand in 1988 with a VAX8700 cpu joining the cluster.

Workstations were beginning to appear; VAXstations with SCSI disks and X-window displays began to find some uses. Using shared network disks under DFS, the DECnet version of NFS, it was easy to share the disks of the main cluster. An Evans and Sutherland PS340 raster display system was acquired to aid the biologists and structure determination of large molecules. In 1989 a DECstation with a MIPS processor was acquired running Ultrix, as a first step using these new and very fast RISC processors.

As a boost to central computing resources, a DEC-alpha AXP-3000 compute-server (Romeo) was introduced running OpenVMS. To reduce maintenance costs this was complemented with an AXP3500 (Alfa) taking over from the mainframe massive 8650 and 8700. These were also configured in a cluster. Some programs were converted from VAX to AXP machine code using DEC utilities (e.g. VEST); the majority were simply recompiled.

Matching increasing disk capacity cartridge tapes with a serpentine layout, TK70 and TK90, were introduced, and the data archive transferred to a CDROM juke-box, reducing need for the 9-track tape systems.

During the long shutdown there was a directive from the steering committee to convert all systems at the ILL to Unix. This is described further in Appendix II. For the scientists HP-Precision RISC systems were given to the nuclear physics and high-resolution instruments, and the other groups were given Silicon Graphics RISC systems. Already in place it had been decided that future data would be in ASCII, based on the TAPDAT extensible format; by compressing data before storage OpenVMS and Unix systems could regenerate files with appropriate read modes (fixed length records/stream_LF respectively)

When the reactor was restarted in 1995 there had been little renewal manpower or money for instruments. Most instruments were still running VAX-VMS. Data were transferred to the central cluster and converted to ascii before being written to the SGI unix server.

The Central Computer Group, Calcul Scientifique

Engineers Siret (Head), De Polignac, Le Sourn
Programmers Tillier, Douieb, Espitalier, Richard, Pinet, Pater
Operators Rey, Metreau, Lagier Rice, Gleed, Airey, Evans

Timeline (approximate dates)

- 1973 installation of rented DEC PDP KI-1070, 250MW 36bit 10 teletype terminals
2x7-(later 9-) track 2x9-track tapes 2xDECtapes 4x40MB disk
Calcomp pen plotter 2xGT40A display terminals
Fortran IV (.F4)
- 1977 acquired. added memory 250MW, additional compatible disk
Tektronix 4010 Graphics terminals (PLOT10)
- 1978 serial line data transfer D17
Versatec raster plotter
- 1979 Replacement Project IBM/CII-Honeywell BULL/ICL/Siemens/DEC started
Benson pen plotter
- 1980 LA36 and LA120 dot-matrix printing terminals
Tektronix 4025 raster graphics terminals
- 1981 Dual operations PDP-KI10, PDP-KL1091S
Fortran77 (.F10)
- 1982 Installation of KL1091S in new computer building
2GB data disks
VT100, PT100 terminals
DECnet, WAN-BITNET (bisync) EARN FRILL
MICOM terminal switch; terminals cabled throughout ILL
- 1984 Macintosh computers introduced, Applewriter PostScript laser printers
- 1986 acquisition of DEC VAX-8600 cluster with VAX11/751 and HSC controller
X25 WAN JANET, EARN FRILL51 on VMS
Pericom Tektronix emulator raster terminals
Ethernet LAN
DEC LN03R PostScript laser printers for graphics
- 1987 upgrade (with ESRF) VAX8650
network tools DEC DFS, DQS, distributed file services/queueing
DEC-Datatrieve for experimental database
Thin wire ethernet LAN cabled to offices
- 1988 Evans and Sutherland PS340 Graphics for Biology (FRODO...)
- 1989 VAX-8700 cluster with VAX-8650 and VAX11/751, PDP10 retired
- 1989 DEC-Ultrix Unix workstation
PGPLOT graphics library
- 1991 Long shutdown
DEC-AXP3000 processor added to VMS-Cluster (Romeo)
- 1992 Unix Conversion
- 1994 Silicon Graphics server, Group HP/Silicon Graphics workstations

DEC-AXP3500 processor (Alfa) added to VMS-Cluster, VAX8650,8700
 withdrawn
 RJ45 cabling of institut with switching routers
 1997 2xDEC/Compaq-DN20 AXP twin processor unix cluster (equivalent to 1
 unit of CISI
 massive cluster
 1999 HPUX 8 twin Xeon processor nodes with Alinka management software for
 multi-processor
 computing
 2000 SUN NFS fileserver introduced
 2006 Fujitsu-Siemens Opteron 8 biprocessor blade cluster (linux) (later
 expanded with infiniband
 interconnect)
 2007 Linux Xeon blades used for central fileserver

DATA FORMATS

With little processing power on the NICOLE and CARINE data acquisition systems initially there was little documented on the data formats on these systems. CARINE data were tailored in contents for each instrument. The Carine Manual (Kaiser et al) described the contents of each field. The NICOLE data were multichannel data blocks, but included a few standard scaler values and a text block which usually held quite a good description of the measurement, but was hand-edited. The data were copied at the end of measurements sequentially to 9-track magnetic tape. In the very early days this was taken to Saclay and written to cards for re-use on the IBM system of the university, but, with the installation of the DEC-PDP10, they were then read each morning on this central computer, and stored on disk. The DECtape data were copied on to the DEC10 by the instrument scientists, and the data were not initially archived systematically, being conserved as separate files.

Archiving Data

The data disk was known as SPCT with a capacity of about 40MB, holding data for each CARINE and NICOLE instrument appended into sequential files for each instrument. Depending on volumes of data these were copied to 800bpi nine track tape. Each run included a ten character field NOMEXP, which was usually changed for each set of measurements, simplifying copying to tape without breaking a user's sequence of measurements. A sequential run number was associated with each measurement (NUMOR). On some instruments (e.g.. D11) this was reset at the start of each year. For other instruments with smaller number of measurements the number was simply incremented.

Data Access

Calcul Scientifique provided a Fortran callable routine to access data either from disk or tape. This routine, MEDOR, written in PDP10 assembler, which allowed return of different number of Fortran fields, depending on the instrument data.. This searched the sequential file of multiple measurements looking for the run requested, like a

good dog. On tape there were often several files, e.g. for the tape HDA, HDA01, HDA02, HDA03 etc., speeding up access by using the hardware skip to end of file tape operation. Although four tape drives were available, the limited disk space for current data (about 40MB) led to quite a demand for access to the data on tape. A major improvement in efficiency followed the addition of directory files on disk indicating the starting block for each run. For tape access too, as the tape was read the start point for each measurement could be retained and bi-directional tape motion optimised. This later routine was known as MEDIR (Ch. Rey).

These archive tapes were transcribed to 1600bpi tapes in 1981, when the old NRZI tape drives were superseded by phase encoded tapes on the new machine. It was also an opportunity to rewrite the magnetic tapes before fading occurred. Subsequent copies were made to optical disk (VAX-VMS) and CDROM. For a small subset of data from D11 which contained the detector data in one file of 64x64x36bits ==> 22480 bytes) the top-end of the spectrum was truncated because the file was copied using the default VAX-VMS block size of 17kB. Fortunately most of these data were recorded in four 1024x36bit records.

When NICOLE was replaced by individual PDP11 16-bit computers the data formats were modified. The original proposal (Hildebrandt) was to record separately the histogram data, the parameters and a text file, and use the RSX version number as the common identifying run number. (Being an octal number was a small complication). A unified, extensible format including this information starting on formal 512byte sector boundaries was adopted (Ghosh ILL82GH05G, ILL89GH3) included these components in a single file with the layout description in the first sector. This was used for PN1, IN4, IN5, IN13, D7, D11, D17, D22 etc. Though 16 bit unsigned integer binary data was initially stored, the time of flight instruments stored a list of overflow addresses at the end. The data were sent by network (DECnet-I) to a concentrator (PDP11/55). Before copying to tape the data were expanded, if necessary, to 32 bit integer data. For D11 and D17 the higher 16 bit data were read and the address and contents of the first 128 non-zero values were stored in the datafile. (Exceeding this number was the consequence of a poorly set up instrument since the detector was likely to be saturated and damaged) The lower 16 bits of data were stored as the histogram. Again on rereading the data for treatment or archiving it was easy to read the overflows before the main histogram and reconstruct in memory the 32 bit data. Given that the instruments could easily produce 200 files per day half of the space was saved and transfer times were usefully reduced. Storage on the PDP10 was always in 36bit integers.

As more instruments were connected to the central computer, standardised header information was included in the start of each data file simplifying archival and access. Initially this was in the form of an ASCII field similar to a Fortran format e.g. 156X 512I 512F 512F 32(1024I). These data were stored on a set of tapes with a series name AAA etc as AAA01 AAA02 etc. Data on the instruments using PDP11 computers were in separate files with names like 12345IN5.don

Most instruments transferred data to the PDP10 using serial communications. To reduce overheads these were polled for activity hence no longer requiring one task per instrument. Typically these were launched either on a timed basis, or manually when data were required urgently. The SANS instruments continued to transfer at

the end of each run.

Under VMS, the data were held on disk in similar files to the instruments though concatenated together. The access on the central system was using DEC-Datatrieve, a relational data base software. Successor routines to MEDIR were provided to allow data access, which also transcribed the 36bit PDP10 data from the existing archive [SPELIB, Blanchard] At this stage the data were identified with the year and cycle of the measurement e.g.. 861, 862, etc. The former PDP10 36bit data were associated with the year/cycle 999. Since the 16 or 32 bit data were essentially in 512 byte blocks, concatenated in the archive file it was convenient to have a single reading function. On the instruments the file was identified by name. From the database Datatrieve routines would find the file (on disk, optical disk etc.) and also give the starting block offset. The same routines could then treat the subsequent data which had the same relative layout.

Conversion to Unix

This was easily achieved with network data transfers; after data had been received on the central system the data were converted to ASCII and compressed. (VMS-COMPILL Ghosh 1992) The TAPDAT file format (described below) was used without change. Using compression allowed data to be retrieved as correct VMS text files or stream format Unix files. Data were simply referred to as separate files using the run number (padded to 6 figures with leading zeros). The hierarchical data files acted to identify the cycle and instrument name (though, of course, this was held internally too). These filesystems were auto-mounted on each unix computers as /usr/illdata/cycle/instrument/numor.Z files. A major advantage of using cross-system compression/decompression was to reduce overheads of storage and network traffic in accessing the data. The counts information once compressed can occupy only ten per cent of the binary data, with no loss of precision. Unless the data read from the network are decompressed onto a local disk the network-load gains are lost. The whole database was progressively rewritten in ASCII files and stored on the SGI-IRIX server, though some errors and omissions required resorting again to the PDP10 tape images saved on CDROM (e.g. the instrument name D11A (1978-79) was unknown to the author of the Datatrieve installation!)

Exporting data

The standardised archive, and the continual pressure for data to be treated away from the ILL to conserve this computing resource lead to standard programs to format the binary archived data to IBM-EBCDIC, and later ASCII text tape files. TRANS (Richard, 1977) included Fortran format statements to read successive data blocks. A simpler approach TAPDAT (Pater, Ghosh 1981) used 80 character records, in a fixed format of text, integer or floating point. Each data field was preceded by an AAAAAAAAAA..., Iiiiiiiiiii..., or FFFFFFFF separator, followed by the number of data entities (and optionally the number of lines of descriptive text preceding the data). The layout was innately flexible and extensible, and subsequently was the basis for the ASCII data used on the Unix systems. The run number (six digits, padded with zeros, e.g. 00klmn) was used as filename since this (or the compressed file 00klmn.Z used) was compatible with all computer file systems, including MS-DOS.

HDF-NeXus binary data

A group of spallation source scientists from Argonne and other centres started discussing a common format for the large volumes of time of flight data that their new machines were creating. The SoftNESS meeting of 1996 centered on use of the Hierarchical Data Format, from the National Center for Supercomputing NCSC at the University at Urbana, Champagne. This had been adopted by NASA and Earth Scientists surveying for their large image datasets. It offered a range of storage and retrieval methods for large files. It was not designed to include much metadata, or MIME additions. Work started on a specification for Neutron Data. Several centres adopted the format, though quite often only partially, older home grown formats still predominated.

Meetings advanced with the adoption of HDF version 4 by several neutron centres. Unfortunately there was great emphasis on raw rather than treated data, hence it was difficult to match the metadata schemes for similar functioning but distinct instruments. If emphasis had been on treated data these differences would have been minimised by the initial treatment. In practice the initial reduction has always been developed locally.

By about 2000 the limitations of HD4 lead to the NCSC development group introducing HDF5. For the first time no backward compatibility was considered; the new design answered the many criticisms of HDF4, with better storage for metadata and better possibilities for including ascii data.

Already a considerable volume of data had been created using HDF4. The neutron data scientists added a complete software layer to HDF to hide (and lose many of the new advantages) of the advances gained in HDF5. It would have been better to draw a line separating the two as NCSC had clearly decided to do. This filtering software gained the name NeXus, for Neutron and X-ray data; by this time synchrotron users were showing interest.

The working group finally coalesced into creating a specification once XML was adopted for providing a prescription of the data and meta-data. This also helped simulation techniques as finally the spectrometer elements could be itemised and assembled into distinct instruments. Each component could be given a relative location on the beam-line, terminating at the detectors.

At the ILL the first instrument with such data was BRISP, a new CRG design with large datasets; this was later followed by other time of flight instruments then the SANS instruments. First D33, with several detectors, then D11 and D22, though there were many illogical differences between the nomenclature and layouts for these last two, despite them having shared identical ascii layouts for fifteen years.

The principal advantage of HDF was that it was often included in commercial data-treatment packages (it was easy to compile and include the library routines in these 4th generation programs like IDL and MATLAB). In practice it was also quite easy to write a few utility routines to read the data with Fortran too! Several centres have worked on generic utility programs for visualising data, with significant investment

arising from the Spallation Source at Oak Ridge being introduced into the more general community.

Whilst these software advantages are significant, especially dealing with very large datasets there remain a number of questions over treating event data, rather than image data. Sparse data are now being measured, which are ill-adapted to the storage methods. While ascii data compressed to a very high level (>90%) the binary HDF data (which includes possibilities for compression) offers a much poorer ratio. Again the storage space for meta-data, while better in HDF5, is not well adapted, for example, for step scanning instruments like 3-axis spectrometers. In fact the XML description for NeXus files offers a useful alternative for including data storage directly, hence sharing the development work of the NeXus group in defining components. These data remain hence in ascii.

NETWORKING

The first steps in networking were linking the PDP11/40 of D17 in 1976 with a serial link which could serve as a terminal or be used to transfer the data using task-task communications programs (Tillier, ILL80T15T). The first function was not used extensively, but the transfers were very reliable. The system was replicated on D11 when a PDP11/40 replaced the NICOLE system. A third instrument transferring data using this ILL protocol was IN6. It was evident that expanding this task-to-task communications feature was quite expensive on the limited central computer resources. For this reason when DECnet was introduced as part of the NICOLE replacement a data concentrator was acquired (PDP11/55) which received data from the ex-NICOLE instruments and continued to write data tapes for transfer to the central computer. With additional computing power the KL1091S received serial line data from many of the neutron beam instruments; a polling program looked for data transfer requests from several instruments. The DEC computer had two additional DN20 front-end PDP11/34s to handle DECNET and other external links. For DECNET phase I only coaxial or synchronous RS422 links were possible; these point to point links were too costly to envisage for general use, and initially were limited to the NICOLE instruments and the PDP11/55 concentrator. D11, using a VAX11/730, had a DMF32 interface which allowed it to be connected to a similar interface on the VAX8600, allowing full DECNET operations. By the time D17 was updated DECNET had been modified to use serial lines, shrinking the cost of point to point connections at 19200 baud. The reduced cost then, and lower cpu overheads of end node rather than routing network licences, was an additional gain.

Initially ethernet was limited to a few instruments, starting with D19, but this was mainly used to offload data treatment to the VAX11/750 associated with the D19 project though the limited cpu capacity of the PDP11/24 control computer could not take full advantage of the increased network bandwidth. The network was extended in-house progressively, and once the VAX8600 was installed thin-wire ethernet was cabled through much of the institute, and DECnet was used quite extensively for remote terminal and file transfer (1987 on). Earlier, indirectly, through the local IN2P3 laboratory, it had been possible to access CERN by DECNET.

Wide-area networking started in the mid 80s with a bi-sync link to the CIGG, University of Grenoble, when the ILL joined EARN, the European Bitnet network (store and forward of mail) with the node name FRILL. When the network was extended through the ILL to EMBL the name changed to FRILL51. The technology of WAN advanced rapidly, and Transpac X-25 access took over the earlier transport layer of bi-sync for EARN, and the ILL joined the French RENATER university X-25 backbone managed by France Telecom. This in turn offered connectivity to the JANET X-25 university system in the UK, and the ILL VMS systems had JANET software installed.

OPERATIONS

The operators, including card-punch assistants, loaded cards for batch jobs, and mounted tapes at users requests from the tape library. An important role was updating the archive each morning with incoming data so that it could be treated. System backup, and disk maintenance too was important. Initially users could use DECtapes for personal copies of important files. Later programs like TAPIP (Le Sourne, 1978) offered 9-track tape storage with the advantage of adding a directory to the contents at the start of the tape. This complemented system backup tools (FSKILL) used to clear off unused files.

System maintenance

The DEC10 system was closed, typically on Wednesday mornings, for routine maintenance. This involved checking the electronics, and the disk head alignments to ensure disk reliability and compatibility of the removable disk-packs (RP03, 40MB, IBM2314 equivalents).

DEC provided a team of engineers on-site to ensure long day system availability. An incipient fault on delivery of the KL1091S lead to the European head of support spending more than 2 days searching for a wiring fault in the backplane using a binocular microscope. Because the ECL logic employed twisted-pair cabling it had been wired-up manually. Towards the end of this marathon a replacement backplane was made available, but the problem was resolved on site.

System developments after 1983

In 1983 DEC abandoned the Jupiter project to build the next generation PDP10. There had been problems of synchronising the various CPU components; the future would hence be an emphasis on advancing the VAX systems. At the time of the acquisition of the KL1091S the performance of the VAX11/780 would have required 3 systems to replace it. By 1983 cluster software was becoming available which would simplify management of such multi-processor systems though clearly the overall power of the single cpu would not be available for intensive computing tasks.

With the acquisition of the DEC-VAX8600 the maintenance contract was renegotiated, allowing greater flexibility for the on-site technicians. The disk subsystems no longer required continued maintenance, and other components were

generally much more reliable. This was a major element in DEC's advantageous pricing for acquiring the new VAX processor. An HSC50 storage controller and a VAX11/751 (an upgrade of the former concentrator, PDP11/55) completed the basic VAXcluster.

The VMS system was very popular with the scientific community; programs could be shared easily with outside groups. The final closure of the DEC10 system was possible when the VAX8700 was added to the cluster.

The CISC instruction sets of mainframe computers required expensive look-ahead logic. On the VAX there were up to four or five locations addressed in each instruction, rendering it difficult to predict and optimise the use of the instruction and data bytes in the cache and number of cycles required in execution.. One example, "find first bit set and clear" for a 32-bit datum was only used by the operating system, but was in the deepest part of the job scheduler. The advent of the simpler RISC architectures lead to huge (apparent) cpu cycle rates. It is however interesting to note that today the Intel i86-64 CISC architecture has developed and maintains dominance commercially. Fortunately the overall power demands have been held in check by the micro-miniaturisation possible with modern nano-scale chip manufacture. Placing large caches within the cpu chip and using other second level cache techniques leave the user oblivious of the real advances in technology using the archaic i-386/pentium etc. instruction sets.

Later there was an addition of an AXP-3000 (romeo) to the cluster in 1990, running OpenVMS with the 21064 alpha RISC processor. This more than doubled the total processing power available. With the final acquisition of an AXP-3500 (alfa) in 1992 the 8600 and 8700 were taken out of service. The HSC50 disk subsystem was replaced by DECstorage towers of SCSI disk packs. One of the AXP systems was kept in working order until 2008, managed by the scientific computing group (RG), allowing the continued use of programs imported from external VMS systems. IP networking had been added to OpenVMS from the start. It also aided recovery of data from CDROMs containing the 36-bit data tape images from the DEC10 era stored in VMS-BACKUP files.

The unix system was completely based on NFS networked file services which shared files amongst the mixture of HP, SGI, AXP-OSF, SUN, and even the AXP-VMS systems. The routine backups remained the responsibility of SI, services informatiques and progressively the system maintenance too for the variety of systems. In the course of time the HP and SGI were phased out, and linux on various PC platforms was adopted. This lead to increased uniformity across the ILL. Scientists still mostly used either Windows or Macintosh for their personal work. Few developed programs, though the gnu packages allowed both Windows and Macintosh to use similar code.

Graphics

Initially there were two GT40 DEC Graphics Displays. These were PDP11/05 computers with a VT11 graphics board accessing the common memory. It offered

scrolling ASCII terminal functions, but the VT11 could interpret autonomously move/draw commands to produce vector graphics on the screen. At the end of the display list a VT11 interrupt instruction allowed the main processor to rearrange the vector list before the display list was restarted. With long display lists the picture started to fade before being refreshed, limiting the total number of vectors displayed. The first general purpose graphics routines were written for this device (using the GT40A library, then BOOK, E Douieb) with the now familiar set of four x by four y axes which became a characteristic of ILL displays. The light-pen was not very useful as a control device for simple plots, mostly being used to control the legendary moon-landing program.

The Calcomp pen plotter was heavily used, but was very slow and not ideal for the many data points coming from ILL instruments.

Advent of the Tektronix 4010 storage tube terminal with attached silvered-paper hardcopy finally gave the scientists an interactive terminal with a useful hardcopy output record. The X-Y cursors allowed selection of peaks which greatly aided data analysis. After two had been installed on the central computer they were also progressively added to the instrument; the hardcopies were shared. In operation the electron beam was positioned and moved on the screen. Secondary electrons were picked off an adjacent grid, and used to rewrite the screen. The two beams were gated to operate independently. Erasure of the screen was by flooding the whole screen and the switching off the rewrite temporarily. The screen had a limited brightness, and was subject to burn-in and fade with use. The basic library PLOT10, (and simpler vector code generators) was available for many systems; it was quite easy to generate the tagged characters which represented the vectors and send them through terminal handlers even on CDC-Cyber systems, and other non byte machines.

In later years the basic PLOT10 MOVABS/DRAWABS etc., commands were interpreted with emulations on raster graphics systems which were much more comfortable to use, though actually lacked the precision of the storage tube with its pure vector graphics. One example was the Hercules graphics board, common on early scientific IBM-PCs. At the ILL the Pericom terminals became a standard, though quite cheap emulator terminals, e.g. the FALCO with a single chip processor, were popular amongst scientists.

A Versatek off-line electrostatic printer took over much of the Calcomp plotting. The plot files were spooled to disk, then copied at intervals to a tape which drove the plotter off-line with its plain paper output.

Colour Graphics

One VS11 16 colour graphics unit was acquired for D19, but proved inflexible to program. For the small angle scattering instruments one ATARI-1040ST program was written (Ghosh, 1985) which took serial ASCII data and displayed one or four 64x64 spectra in 16 colours (display using 600x200 pixels) This ATARI microcomputer had a graphics chip with a colour lookup table, making possible a wide range of colour representations.

Over 10 ATARIs with associated SONY display monitors were connected by serial lines both on instruments and the central computer. The display driving programs generating the ascii serial data could use a shared memory segment which was available both on the instrument computers and the central VAX computers. Some were used for monitoring incoming data, others for treated data on D11, D17 and D16. 35mm cameras could be attached to screen-mounts and using polaroid colour films these yielded transparencies for subsequent projection of results at modest cost.

Benson colour pen plotters were also much used, especially for crystallographic structures and gamma-ray level schemes.

Graphics software was primarily based on PLOT10 and the pen-plotter standard libraries. Computer Associates DISSPLA had some usage, but programs were not transportable.

From 1988 VAXstations with graphics capabilities were installed on instruments. On D7 much use was made of the VAXstation-II GPX graphics (Schaerpf) (pre X-window) using Basic. The X-window graphics were not easy to program, lacking documentation. At this stage a commercial package by PVwave was acquired, and this led to a rapid development of display software using the IDL language. Later as PVwave showed less interest in scientific display contracts were established with IDL, who had meanwhile developed their own graphics capabilities for these systems. The cost of licenses remains a significant recurrent cost, though it is possible to create exportable programs, (LAMP, Richard, COSMOS, Cubitt, etc), containing a free run-time IDL kernel.

For displaying large biological molecules an Evans and Sutherland PS340 display system was acquired (1988) and used primarily by the biologists. This model, with a raster display ran autonomously from the PS2 vector display which was usually integrated into a VAX system. Typical programs used included FRODO for rearranging parts of molecular strands.

PostScript

Graphical output as PostScript files became increasingly common after the introduction of general networked printers, apart from the Macintosh appletalk network. Poster printing and embedding EPS files in documents became common, helped by the initial acquisition (with contributions from the ESRF and EMBL) of an HP PostScript A0 printer. The cost of this type of printer rapidly dropped, and PC/Macintosh software (Adobe Illustrator, PowerPoint etc) evolved to allow easy creation of posters of scientific results.

Terminals

Initially the DEC10 console computer was a KSR35 heavy duty teletype (10cps), though this was soon changed for a DECwriter LA30 80 column matrix printer (30cps). A VT05 VDU terminal offered dynamic visualisation for controlling processes. The users initially used 6-8 ASR33 teletype (electromechanical 80

character 10cps cylindrical print-head) and two GT40 displays. The teletypes were located around the ILL, each connected with a current loop circuit.

The introduction of two Tektronix storage scope graphics terminals (13 inch) and their 4610 hardcopy (1978) led to rapid development of interactive graphics for viewing and analysing data. Line plots, complemented by 2D contour plots of multi-detector data. The Diffraction group actually financed one of these terminals (Hewat) and had priority of use during part of the day. Refinement of powder diffraction data was a production activity (Hewat). Because true vectors were drawn the diffraction patterns were of much better quality than those later produced on raster screens.

From 1978 the DECwriter II, the LA36 9 pin 132 column printing terminal (30cps) was being introduced. Most were subsequently upgraded with a third party electronics board (FUNGUS) which could print at 120 cps, with bi-directional optimisation. Though not capable of true descender characters it could print upper and lower case.

When the DEC1091S was acquired additional DEC VT100 vdu terminals, and later Plessey PT100 compatible vdus were added to the park. A small number of Tektronix 4025 raster (640x350) graphics terminals were also acquired, though the resolution was inferior to the storage screen terminals; the scrolling alphanumeric capabilities and bright displays were quite welcome. We discovered that the operating system was actually in upper and lowercase for the first time; terminals had full character sets, and the original LP10 drum printers were complemented by two 1200 lpm LP20 charaband printers with full characters sets.

The formatting program RUNOFF was used for documentation with a daisy-wheel printer for improved quality. DEC later produced the DECwriter III, the LA120, which printed at 120 cps, and became standard as console printers on instruments. The VAX11/730 computers were delivered including LA100 console terminals (120cps, desktop 120 character wide printers).

The Pericom Grafpak (13 inch 1024x800), then the Pericom Monterey MG200 (17 inch, 1986) were the first modest cost raster graphics terminals which had widespread use. Later the much cheaper single chip Falco (13 inch) was cheap enough for scientists to have graphics terminals on their desks. Matrix printers could print screen copies. All these terminals used the Tektronix PLOT10 ASCII print coded graphics, connected via serial lines (9600 baud).

The terminal network was extended progressively throughout the ILL from 1982 with the introduction of a MICOM, (later Satelcom) terminal switch; this allowed access from any terminal to either the central system, or the instruments. One terminal on each instrument also could use this to access the central computer.

Following the introduction of unix, X-window terminals were added to the network; there was less need to augment the number of workstations. The cost of the 19-20" CRT colour monitors, the principal component dropped rapidly. The advantages of a non force ventilated unit in offices was appreciated. This again changed with the

advent of flat screen TFT monitors which liberated considerable desk-space taken by the CRT screens.

The X-window server in Cygwin PC software and the gcc compilers enabled building versions of unix software to run on PC-Windows. The large LAUE package from John Campbell at Daresbury (40000 lines of Fortran 75000 lines of C) were slightly modified for use on Windows. Some scientists used commercial x window servers on PCs (Xserve). With OS X, Macintosh computers had XQuartz X window server software built-in for using the unix graphics.

HIGH-PERFORMANCE COMPUTING

Instrument simulation and ab initio calculations required increasing performance only obtainable by distributing processing in parallel over multiple systems. Early work was done on remote systems accessed over the network. One system was the CISI Cray T3 computer system incorporating 256 DEC-Alpha processor nodes. At the ILL the first steps in cluster operations involved installation in 1996 of two twin processor DEC-DS20E alpha-EV6-833MHz systems, with fast memory interconnect. These operated as a module matching the CISI Cray T3 basic four processor unit. The system proved to be very powerful, and only when the much faster 3GHz Xeon cpus were more common that they could be retired. The first linux cluster at the ILL used HP bi-processor computers, essentially large PCs. Portland (PGI) Fortran was chosen, partly because it had been used by software developers, and partly it offered a complete package, including BLAS routines, with ready built parallel processing extensions. There was often a mismatch between software component versions for other suppliers. It was still necessary to build some of the low-level BLAS routines on these systems. A management tool, Alinka, was also acquired, to manage the batch use of the set of processors. Later installation of a group of AMD computers, (which had superior shared memory hardware to the Intel processors) indicated that a small number of users could mutually agree to share a system reasonably equitably without the constraints imposed by Alinka. Initially gigabyte ethernet linked cluster elements. When Linux blade systems were introduced Infiniband interfaces were later added to improve system interconnectivity; most programs used the MPICH task synchronisation package..

PERSONAL COMPUTERS

As the price of personal computers dropped scientists acquired either PC-WINDOWS or Macintosh systems, with some groups have preferences for one or the other. Once thin wire ethernet had been installed the appletalk network was replaced and all systems used first this then the RJ45 cabled switched ethernet network. The number of workstations too increased. After the initial period (1990) when a DEC-Station (MIPS-Ultrix) workstation had been acquired, from 1992 HP-UX (Nuclear Physics, TOFHR, 3-axis) and SGI (LSS, Diffraction) were preponderant, The OpenGL graphics of SGI was not heavily exploited and though the SGI system was the most comfortable in use the company finally exited from the field, and linux became increasing used at the ILL, though only a few scientists used it for personal work. The central data files were moved from an SGI server first to a SUN server, then in the early 00's to linux based systems (SuSe).

SOFTWARE

The principal language at the ILL was Fortran. This was initially Fortran-IV. The instrument control system on CARINE was written in Fortran with a real-time layer (Le Sourné 1972). On the central computer Cobol was used for an administration activity, management of mission payments. The data archiving programs were written in PDP10 assembler. The Fortran-10 on the DEC-1091 was a superset of Fortran-77. The VAX offered the features of language independent compiled object code which allowed mixing languages.

Mathematical Support

The two applied mathematicians, Siret and de Polignac were supportive of new techniques such as introducing Bayesian analysis (with Mendes in initial years). This was supplemented by encouragement from Hossfeld (Juelich). Quantitative Maximum Entropy (from the Cambridge Group with Daniell and Skilling) was presented by Sivia from RAL who made a number of visits. For the most part scientists used conventional least-squares fitting procedures for model fitting which also provided some indications of errors in the final results.

The reactor division performed calculations using standard packages on the Central Computer for source design (Ageron, Robert), and several large international standard packages were installed.

Numerical support (Siret, De Polignac) maintained use of well-known scientific software libraries, the IBM Scientific Software Package, LINpack and CERN libraries, the Harwell libraries and the UK NAG libraries etc. For the PDP10 the source code was usually compiled and tested in-house; for the VAX computers pre-compiled libraries were available. Later for Unix most libraries were recompiled for the four common unix systems, HP-UX, SGI-IRIX, DEC-Alpha, and linux. With linux came the possibility of license-free GNU software, though for high performance, multiprocessor Fortran, a PGFortran license was bought., which included multi-processing support. The most expensive licenses were the recurrent costs for the IDL package which formed the basis of several instrument display and treatment packages (e.g. LAMP, Richard)

Standard ILL utilities for accessing data were maintained by Calcul Scientifique Initially written in assembler these were mostly written in Fortran. The VMS system had a common object file allowing languages to be mixed. Later linux systems using gcc had interoperable language capabilities, and with but minimal changes the same programs could be recompiled to run on PC-Windows, Macintosh, and Linux, using appropriate versions of the PGplot graphics library, IDL, Matlab etc.

INSTRUMENT CONTROL AND ELECTRONICS

The control of neutron scattering instruments required precise positioning and fast counting electronics, together with efficient detectors. The treatment of data in user institutes was not common; the visiting experimenters were expected to take their data away. Given the unique instrumentation there was however pressure to provide means for data inspection during the measurements, to ensure efficient use of instruments.

The cost of instruments was initially divided equally between mechanical construction, electronics, and control computers. With such expense the computing support groups were staffed with a number of engineers and programmers. For subsequent data analysis the scientists were primarily responsible; the central facility provided staff to operate and maintain a data archive, and offer some technical support, numerical and graphical libraries, for the scientists, but this was quite limited on the instrument systems.

At the ILL computer instrumentation and electronics was initially headed by Axmann; after his return to Germany in 1975 the responsibility passed to Drexel.

Timeline

- 1972 CARINE (Telemecanique T2000), NICOLE(Telefunken TR86) installed
- 1975 DEC-PDP11/20 (RSTS) for GAMS, BILL, DataGeneral Nova IN7
PDP11/40 (RSTS) D3
PDP8E D15, D16
- 1975 PDP11/40 (RSTS) IN10
- 1976 PDP11/20 (RT11) IN11, PDP11/40 (RSX11D) D17
CAMAC Memory system
- 1978 PDP11/40 (RSX11D) D11A, D11B
The deuxieme souffle and connecting to a new central computer
- (Appendix I)
- 1979 PdP11/40 (RSX11M) D8, D9
- 1979 PDP11/34 (RSX11M) IN4, IN5, D7, PN1
SEMS-Solar 16/40 D1B, IN12
MuMu microprocessor axis control (Munnier)
- 1980 PDP11//34 (RSX11M) D8, D9
serial line microprocessor counting, positioning (Falaise)
Solar 16/40 IN1, IN3, IN8
- 1981 PDP11/34 (RT11) IN6
DEC-CAMAC controller
- 1982 PDP11/34 D19
- 1982 PDP11/23 (RSX11M) D1A, D2
- 1983 PDP11/24+VAX11/750 D19
- 1984 PDP11/24 (TSX) IN20
VAX11/730 D11
- 1985 PDP11/73 (TSX) IN1, IN3, IN8
VAX11/730 D17
- 1986 vax3500 IN4, IN5
VME electronics for counting and chopper control

	vaxstation IN6, D7
1988	microvax-II PN1+Labview microvax-II D19, D9, etc
1989	microvax-II D11, D17, with vaxstation graphics, VME, PC ancillary control
1991	Long shutdown
1992	Unix transformation planning see appendix II

Axmann head engineer to 1975 (initially Julich, subsequently HMI, Berlin)
 Taeschner assistant engineer to 1975

Engineers Barthelemy, Schwab, Grevaz, Kaiser, Ledebt, Messoumian, Descontes
 Programmers Hildebrandt, Blanchard, Roure
 Technicians Tschofen, Pastor, Delacroix, Guellec

The CARINE system at the ILL IN1,IN2,IN3,IN8,D1A,D1B,D2,D4,D5,D6,D8,D9,D10

Engineers Kaiser (3-axis), Barthelemy (Diffraction)
 Programmers Roure, Blanchard
 Technicians Pastor, Delacroix

The Telemecanique T2000 process-control computer used for the step-scanning instruments, CARINE, was novel in that the multiple instruments were served by swapping separate processes in and out of memory from disk on demand, typically at the end of each counting step. There were two operational computers, with a third for backup and program development. Each system could run up to 6 instruments with an additional background process. The computers had a 19 bit word length, matching a twelve-bit positioning component and a seven bit operational code. Each process could have a maximum of 32k words, and were swapped in and out from a magnetic drum.

The user's terminal on the instrument was an ASR33 Teletype (printing at 10 characters per second); it was possible to set up a sequence of commands for execution. An additional button labelled APPEL requested the computer, interrupting the current activity.

To offer maximum flexibility a semi-compiled FORTRAN had been developed (Le Sourné) and this was heavily used, especially on the 3-axis instruments, to set up scans in reciprocal space and energy. Typical counting times at each position was in minutes, hence the activity demands were limited, though the alignment times involved very short counts.

For diffraction, especially the four circle instruments, the counting times were very short and a huge fraction of measurement time was lost due to the slow interaction with the control system. This was a major factor in the advantages of updating the computers by more modern systems.

The counting electronics were based on single detectors, monitors and scalers, with time and monitor scalers available to control scans. The motors were controlled

through the standard ILL control box using Precilec coders multiplexed to a single decoding electronics.

In addition to running the instruments one additional background process could swap and share the cpu; this offered development facilities using Fortran (and some graphics applications too were available e.g. for plotting D1B powder data). Once the PDP10 was available there was little development of such applications given the priority taken by the instruments.

The NICOLE system at the ILL IN4,IN5,D7,D11,PN1

Engineers Kaiser, Grevaz, Schwab
Programmers Hildebrandt, Wotschack
Technicians Anderlohr, Tschofen

For other instruments multichannel data acquisition was required:

- Time of flight spectrometers IN4, IN5 (inelastic, quasi-elastic scattering)
- Diffuse scattering D7
- Small angle instrument D11
- Fission fragment spectrometer PN1 (Lohengrin)

The principal function was starting, stopping counting and data storage, with limited data visualisation of the histogrammed data.

These were performed by a single AEG-Telefunken TR86 computer; the system was called NICOLE. Elsewhere at this time in Europe and America dedicated instrument control computers were being introduced to offer greater reliability and flexibility.

History has it that the system kernel was conceived at the university of Karlsruhe as a theoretical problem of real-time computing. I found the following description of Prof Herbstreith below:

Heinrich Herbstreith

After graduating in Communications Engineering at the Technical University of Karlsruhe (KIT today) I began my research in IAC (Institute of Applied Nuclear Physics), headed by Prof. Beckurts, who was later murdered by the RAF. He founded the working group under the IDT physicist Dr. Krüger later dean of the faculty of computer science of the University of Karlsruhe. The group consisted of eight graduates of physics, communications engineering and mathematicians, all of whom had already knowledge in programming. The first project was to develop an automation system using computer science systems in hardware and software. The name was CALAS (computer assisted Laboratory Autotomation System). We developed hardware components for data acquisition and graphical output (screens) as well as the real-time software (Real Time Operating System), the data acquisition and storage as well as the graphical output of the neutron spectra. The target computer TR86 Telefunken from Konstanz was chosen. It was a process computer with real-time interface for the connection of data acquisition devices with 64kb memory, hard disks, magnetic tapes, punched card input and output, as well as

graphic screens. The system was put into operation after two years for eight experiments and after one year it ran smoothly.

The second project was the implementation of a similar system called NICOLE for the RHF in Grenoble, whose director was the Nobel laureate Prof. Mössbauer. I was the project manager and after two years the operations were satisfactory. In 1972 I was appointed professor to establish a department of computer science at the University of Applied Sciences Karlsruhe, In 1973 I was at 32 years the youngest professor in Germany.

The computer engineer for NICOLE was Taeschner with Schwab, technicians Anderlohr and Tschofen, and with Wotschack as programmer in charge of PN1 acquisition and Grevaz and Messoumian for D11, IN4 IN5 and D7, with Hildebrandt as programmer. Programming was essentially in machine code. The interface electronics was under the responsibility of Klesse, and Ertel looked after the CAMAC modules, notably the time of flight encoders. These were derived from accelerator physics designs, and consequently had a finite reset time, corresponding to awaiting the next valid accelerator pulse. For the rotating choppers at the ILL where there was never any uncertainty in the next trigger pulse from a rotor at 2000-20000 rpm and there was always the inconvenience of the incomplete recording of spectra with this time-gap. The time generation was set using the computer, or via thumbwheels; the units were half microseconds, (50 microseconds for D11) using octal (though decimal thumbwheels were used).

The computer was a 24 bit computer with two Burroughs disks ("Boris" and "Rebecca") each with a fixed vertical platter about 1m diameter, driven with a belt drive, with 800 bpi magnetic tapes, card input and output, and a line printer. A second machine served for development work. Initially, before the Central Computer was installed, data tapes were sent to Saclay, being returned with a copy on cards. My first experience of this service unfortunately coincided with the tape being shredded at Saclay. Data processing at the ILL was then provided through an RJE link (the "Terminale Lourde", printer and card reader) to the (IMAG) University's IBM 360-VM computer. Later the tapes were taken off each morning and read on the PDP10 directly. Again unfortunate data losses occurred in the early days when the previous days tape was replaced again on the Nicole tape drive and overwritten.

Each instrument was connected through several video-grade cables. These carried the event data (as parallel descriptors up to 16 bits) which were sent through a FIFO buffer on the instrument with a handshake on the computer. For D11, with count rates from the area detector which could exceed 100,000 events per second, data were accumulated in memory (4096 words). For the other instruments the mean data rate was much lower, and the data were updated periodically rewriting the incremented data back onto disk. When computer problems arose these data were safe, but the current data from the small angle interment were lost on restart. Fortunately these latter experiments were typically short measurements (less than 30 minutes), and the experimenters manned the instrument continuously, and hence

were present to restart the computer. While D11 had a maximum count limit of $2^{24}-1$, it was only discovered in the last months of operation that the limit for the data stored on disk for the time of flight instruments was $2^{20}-1$. The less demanding time of flight experiments never attained this limit until an experiment to measure the scattering from highly absorbing liquid ^3He in 1975. To observe the small signal counting was continuous for over two weeks. The beam monitor data exceeded this count limit several times and posed major problems in normalising the data subsequently.

The User Interface

The premiss for the design seems to have been built on a shared multichannel analyser. As a consequence the user was presented with a display (HP) CRT monitor, a keyboard, and a console with rotary switches to control x,y and z ranges, and a set of control buttons, together with a cursor intensity display (on most units in binary, later on one or two units in decimal). The buttons selected x,y preferences, and an isometric display (useful for the 2D detector on D11). The TR86 monitored the box and responded with the appropriate display. This required the counting electronics, the console box, and its switches, and the monitor to be directly connected to the computer. Several 3cm diameter multi-core cables connected each instrument to the controller. These were later used (without permission) as a general support for later cable installations around the reactor and guide hall. These high grade cables terminated with LEMO connectors with around 100 pins. Many years later a number of these connectors in pristine condition were still in the spare parts stores of the electronics group.

The screen was divided into three regions: a graphical or text display region, a keyboard input line, and a line for error messages at the bottom.

In graphics mode the data from a scratch buffer was displayed in various forms and the scale could be change using rotary knobs on the console which were sensed by the processor. In text mode the text could be edited. This was used in the data save routines; the first twelve or so values were read from electronics (beam monitor, time etc.) and the remaining lines were comments. This was used from 1974 onwards to annotate the data with information not directly readable from the instrument, but which was of great use in later treatment of the count information. A cursor (haeltespunkte) could be set with thumbwheels, and the intensity content displayed (in binary).

Although CARINE had a silvered paper hardcopy in the machine room which was used often for D1B few users of NICOLE made hardcopies of the screen there. Data tapes were changed each morning and copied to the Central Computer (a DEC-PDP10).

The keyboards had the German layout. This hence added variety to the American keyboards on the ASR33 teletypes used on CARINE, and the IBM029 card punches with French keyboards.

There were a number of small programs available for the user. Obviously the basic start/stop commands could be sent; the preset time or beam monitor values were set by thumbwheels on the CAMAC or sent as commands.

A standard program, PAPE, could transfer data from disk to memory buffer for display, or to tape. (Command STA,PAPE,BILD or STA,PAPE,BAND), or bring the text parameter block into memory for on-screen edition of supplementary information. About 64k words were available to buffer data for plotting in several histogram modes controlled by key switches.

The principal control for all instruments was the preset time or beam monitor limited count. D7 had a number of motors too which could be set, while PN1 could control magnetic and electric fields used in the fission fragment separator. On D11 even at this stage a number of timed rapid stop-flow experiments were satisfactorily performed.

Program development was performed on a second machine which was also used as a reserve of spare parts.

The overall reliability of the NICOLE system was low, and successful operations depended to a large extent on D11 experimenters, who were present at the instrument night and day, to restart the system. The current D11 data (in memory) was lost, but the other machines conserved their data on the disks. The real-time system made it inherently difficult to reproduce problems. In one incident, in the twilight days of operation an error message appeared for the first time, leading to great celebrations in the machine room.

While D11 used 24 bit integer storage, during one of the last experiments on IN5 before the system was abandoned, an experiment ran for over a week (measuring the minute scattering from ^3He). Normalising the data subsequently showed that the incrementation only treated 21 bits before "overflowing".

The NICOLE replacement was a project of the Deuxieme Souffle, with its replacement using individual PDP11-34 computers on each instrument being completed in 1978. As instruments were successively transferred the overall reliability inevitably improved.

Other computer systems

IN7, a statistical chopper, had a Data General Nova 16bit computer. The nuclear physics instruments, the Gamma spectrometers Gams1, etc and the Electron spectrometer Bill had simple PDP11/20 systems, delivering data on DECtapes which could be reread on the central computer. These used the RTS system (Ledebt, ILL) developed in-house initially, before DEC produced the higher performance RT11 systems which replaced it. RTS was used too, initially, for the PDP11 on the D3 polarised neutron diffractometer (a modified Harwell Mark VI diffractometer). D15 and D16, Harwell Mark VI diffractometers were originally under construction for use in the university program at Harwell, but on Britain joining the ILL it was decided to install them at the ILL, together with their PDP8E computer running the Harwell ANDROMACHE software, further developed at the ILL by the loaned Harwell scientist Stuart Wilson. Data were transferred to the central computer on DECtapes.

Networking

The first instrument to be connected with a serial line to the central computer was the PDP11/40 of D17 from the start in 1977. The protocol allowed use of the instrument DECgraphic11 display as a GT40 terminal of the PDP10; more important it served for data transfer using a check-summed protocol (Le Sourne, Tillier).

PDP11 support

Tschofen was trained in supporting the PDP11/40 hardware. He demonstrated his aptitude when the computer on D11, as always under extreme pressure from the heavy scientific program imposed on it, failed one morning. With Metrix voltmeter in hand he programmed test sequences which he keyed in to the console, then following the logic sequence on the cpu board, identifying a misbehaving gate chip. After soldering in a replacement the measurements could continue before lunch. (Some years later he identified a fault on the PDP11/55 concentrator backplane wiring which had always eluded the DEC technicians)

For software, Roure was sent for training on the RSX11D system, and wrote the low-level handler for the CAMAC crate controllers. This included full multi-user protection to forbid interactions between the crates used by D11A (SANS, Ibel) and D11B, (Diffuse scattering detectors, Kistorz). Later, under RSX11M, there was considerable simplification and these distinctions were rendered unnecessary.

Electronics at the ILL

Engineers: Klesse, Lefevre, Falaise, Thillosen

Technicians: Munnier Schuba, Epaud, Taffut, Vernier, Bunaleau, Valentini, Bresson

NIM modules were used for rate meters, and as a powered chassis for detector amplifiers etc.

CAMAC was used extensively as the primary interface to computers.

The CAMAC at the ILL was quite unique. Elsewhere it was used only as a computer interface; exceptionally LEDs signalled a limited functionality. At the ILL this was developed into a mixture of control and display. Many units had switches and buttons, only some of which could be overruled by computer control. This allowed some off-line control to be performed, performing start/count/stop operations. In addition the scalars had either an integrated display, or a generic display which operated by stealing from the CAMAC bus cycles to interrogate units during latent periods. Preset timers used thumbwheels to set values (two digits, $KL \times 10^{**M}$) - these bcd scalars could hence not be reloaded precisely to continue interrupted sequences). A number of modules had reset buttons. There were no serial numbers on the modules, which often were permanently wired in different configurations rather than having internal jumpers. There was little written documentation, apart from the original circuit diagrams. The French engineers sent out designs for external constructors. Klesse had additional technicians to build boards in-house.

Axis motors and control

A generic controller for motors had been developed to use the Precilec coder and reader. This used a coarse and fine coupled coils to measure their relative positions using high frequency AC. The input was switched using a combination of a rotary selector and keys; only one shaft could be read at a time. As a consequence position setting was performed sequentially. This too was reflected in the initial mechanical designs using rubber wheels pressed onto the floor by air pressure, and this too being used to brake static components.

Instrument setup at the start of each cycle was very slow using computer controlled short scans, as the computer had to service requests from multiple instruments sequentially. Many scientists were very adept at using the manual positioning keys and scaler start-stop buttons to avoid the delays in awaiting the computers control.

Motors were typically step motors or DC servo motors for the axes with large inertia. Later large step motors were employed even for major components e.g. IN8.

Some instruments (notably the British contributions built by Grubb-Parsons, D3, D15, D16) used absolute encoders which enabled simultaneous positioning. These were however more fragile than the Precilec coders which could be used even in high radiation areas notably the monochromator casemates.

Counting systems

The multichannel instruments passed data along parallel links to the original computer systems. The remainder used CAMAC scalers.

Deuxieme Souffle

The replacement of the Nicole system by individual PDP11/34 computers included new CAMAC counting electronics first introduced for D17; an incrementation unit took in data and updated memory modules using a separate front data bus cable. Data were read out using an autonomous unit which placed data on the CAMAC bus for DMA access to the PDP11 memory with block transfers. The CAMAC configurations exceeded a single crate, and CAMAC Branch configurations were used. The memory system (Klesse, Axmann ILL76KL336, Klesse, Mounier ILL77MU18T) had three write cycles and one read-out, though at high data rates the latter was replaced by an additional increment-write. 16 bit registers were used to control input and output ranges. For IN6 this was extended to 18bits using four word alignment to cope with the larger detector array. The memory units included first 4k, then 32k, then 128k 16 bit words which could be used concatenated as 32 bit storage. A microprocessor display controller with a plasma display screen offered visualisation of the data being acquired (Dimper). With Tektronix displays on the PDP11s more sophisticated display programs were developed, common to RT11 and RSX11 (Knowles, Nelson). IN6 was the first instrument to use the DEC CAMAC and DMA controllers which connected it to the UNIBUS. For later multi-crate systems each crate had a separate controller.

The NICOLE instruments ran the RSX11/M operating system with DECnet version 1, which allowed them to transfer data to a PDP11/55 "concentrator" which had considerable intermediate storage using DMC11 interfaces and pairs of co-axial cables. One of the functions of this machine was to write the data daily to tape for

transfer to the central PDP10. This also involved modifying the data to 32-bit storage if overflows had occurred; on the instrument data these events were stored after the end of the detector array data.

The full replacement of CARINE was a more piecemeal activity started later, though PDP11/40 computers were already operational on D8 and D9.

Although the Solar 16/40 computer had a better performance, the software and disk layouts were rudimentary and inflexible. By 1980 the DEC PDP11/23 had been produced, a single board computer, with commensurately lower price, using a multiplexed Q-bus rather than the DEC-UNIBUS. This could be used in conjunction with the Schlumberger CAMAC crate controller. Simple instruments like D1A used the LSI-11 together with a 2.5MB RK05 cartridge/2.5MB fixed disk (same spindle) running RT11. The RK05 cartridge was a common feature until replaced by the RL02, 10MB cartridge (pre-formatted). Instruments had an unlimited autonomy, and the set of data cartridges could be re-cycled. Data were initially copied to tape on the PDP11/55 concentrator for transfer to the central computer.

The typical instrument configuration included:

PDP11/34, 256kB memory, 4 DL11 serial i/o connections 2.5MB RK05 cartridge/2.5 MB fixed disk

DEC-CA11F Camac crate controller (+CA11FP DEC-Camac DMA controller for multi-detector memory readout)

or PDP11/23, 256kB memory DQ11 4 line serial i/o and Quiverter bus adapter for Camac

Introduction of 32-bit computers

By mid-1984 most instruments were being controlled by PDP11 systems. D19 the large multi detector diffractometer was controlled by a PDP11/24 (essentially the LSI-11/23 computer with a UNIBUS backplane), connected to a VAX11/750 for data treatment. This enabled the same positioning software to be shared from the other diffractometers (Allibon). The constraints of treating large arrays of data on the 16-bit PDP11s had always required some skill in program design, and usually required overlaying programs to bring the code down to the 64 kB limit of the 16 bit address for a single task. For RT11 the link operation was very rapid, but for RSX11 the task-builder, TKB, was very inefficient, building each overlay separately, often requiring 30-40 minutes to construct a task. There was hence much more rapid development of control and data treatment software on the RT11 systems. The initial constraint of the RT11 foreground/background system was circumvented when TSX11 was introduced. This was a proprietary multi-user version of RT11, though never included DECnet compatibility at the ILL.

The VAX11/730 was available at this time, with a UNIBUS which could control CAMAC directly. With initial help from an RAL engineer (Pulford) who had been involved in linking their data acquisition systems to the 11/730, a CAMAC handler was written specific to the DEC CAMAC controllers (Rice,Liles). These did not conform to normal VMS standards, since both serial and DMA control operations were possible. A VAX11/730 was installed on the instrument D11 in 1985; there were notably few problems with the low level software. Up to four or five programmers

and users could be supported by the system (there were no graphics displays, other than Tektronix terminals. Program linking was rapid. A distributed approach was made to the control program. Each motor had a separate control task. The memory read was a separate task etc. Some 24 tasks were initiated and sent into a hibernating state when the system was started. A master scheduling program sequenced operations, waking up desired activities. All communications was through shared memory, hence there were never problems of file and data formatting. a monitor program too could show any parameter in the shared memory; this would write to the first 6 lines of a VT220 vdu terminal, while the rest of the terminal display scrolled, and took users' commands. VMS allowed unrestricted program sizes, and rapid development of the control and locally available data treatment followed, together with the inclusion of the ATARI colour displays. When D17 was similarly equipped the control program was identical, but the characteristics of the instrument followed configuration of a different set of motors etc in the initial start-up procedure. The system was tested on one afternoon, simply moving the UNIBUS cable from the PDP11 to the VAX, running without any problems. It was introduced without further delay. When later adapted as a reflectometer it was easy to add new motors. Extra copies of the motor task were added in the startup phase. In the shared memory there was space for up to 64 such subtasks; a snapshot of the shared memory served in restarting the system, holding the instrument identity. It was never necessary to write a formal management tool for this space; all software contributors respected a single page table describing the known sub-tasks and who could write specific values.

The initial configuration included 2MB of memory. This was expanded to the 5MB limit. An IEEE instrumentation interface too was added. This was connected to a set of digital voltmeters through an optically isolated serial link, protecting the computer against mishap. A pool of Keighley DVMs could be shared around the group. They were controlled and read using ASCII control sequences. There was space reserved in the data files for up to 8 values taken at the beginning and end of measurements.

32 bit computers hence became attractive for instrument control. The next systems obtained were VAX3500 computers with a Q22-Bus (IN5 etc using VME). For other computers a cheaper solution was to buy second-hand microVax-II computers, with Q22-buses, and use QUNIVERTER bus conversion cards to attach to the UNIBUS CAMAC controllers. Only one software modification was required in the handler - it was necessary to extend the timing delays to allow for the higher processor speeds compared to the 11/730. The microVAX-II became a common configuration for ILL instruments towards the end of 1990. As VAXstations dropped in price D11 and D17 added these as graphics devices, with a utility program to maintain a copy of the control computer's shared memory on the workstation. Using the shared file-system (DFS) the users were never aware of the two distinct systems. For the nuclear physics on PN1 LABView was introduced for flexible control.

Other computer control systems

While not formally dependent on the ILL, the special projects often shared concepts. The Neutron Dipole moment experiment (1980 to the present, Sussex-RAL collaboration) initially used a GEC4080 computer, connected through serial links to

BBC microcomputers. In 1985 this was changed to a VAX11/730. As the ILL staff had found, there was an immediate gain in ease of use and productivity with the VMS system.

Early microprocessor projects

Introduction of the Solar computers required some developments to link to CAMAC for IN12. A special parallel interface was developed using the Intel 8080 (Ledebt, Lefevre) to create a message box in memory. A CAMAC based microprocessor system using the Intel 8085 chip was developed for positioning axes in 1977. (Klesse, ILL79KL5T, ILL79KL6T, ILL79KL7T).

Later for the other three-axis instruments serial line modules were developed to control positioning and counting, (Falaise ILL84FA3T, ILL80FA39T). A serial line multiplexor linked these modules to the Solar control computers as a way around the lack of directly linking to the Solar computer bus.

In 1990, in collaboration with hardware from RAL, an INMOS transputer network to decode multi detector data was simulated successfully, but the flash ADC hardware was never completed at RAL. The system was hosted initially on the VAX11/730, then, much more efficiently, using an SCSI interface on a VAXstation. Up-scalable data rates were achieved starting at 30000 events/second.

VME-electronics

By the late 1980s there were increasing problems in maintaining CAMAC modules (even finding replacement chips), and the distributed processing possibilities of a chassis of VME modules was attractive, offering, for example, phasing control for multi-choppers, together with control of counting. In practice the demands of such instruments were much simpler, and a master/slave organisation was finally adopted, similar to the CAMAC systems. Newer and cheaper instrument-bus systems have been developed, using smaller profile boards. Fortunately the option to develop even more expensive VXI applications (with an analogue part of a larger format board) were not adopted.

The VME boards used Motorola 68000 chipsets. These were interfaced and developed using Macintosh computers. Initially at the ILL, CERN software was used to link to the VAX instrument computers. This was later changed to OS-9 then further developed using a configuration database when the VME systems were connected via Ethernet on a dedicated physical network. Considerable use was made adding ILL designed FPLA custom daughter boards to pre-manufactured VME mother boards.

Development of arrays of 1D position sensitive detectors lead to collaborations between the detector group and the electronics to use flash ADC converters for localising positions. This started the present tendency to deliver interpolated data to the scientists. The actual derivation of the conversion to real space is not included in the current data.

Elsewhere at the ILL the sample environment group used 6809 chips, running OS-9

for the ILL precision sample temperature controller, which was connected using a serial link with check-summed communications (Dagleish).

Other Computing activities

Office automation started with secretaries using VT100/PT100 vdu based word processing running on a PDP11 RSTS system (1981, Grevaz). Work for the scientific programme, scheduling and the regular meetings of the scientific subcommittees were also developed in-house. From 1985 PCs/Macintoshes were introduced and Microsoft Office became the standard software for word processing and presentation activities. Later (1999) HR projects were developed in-house using Oracle on HP system hardware which were housed in part of the computer room no longer occupied by the behemoth main-frame systems of the 1980's.

The ILL web server was started in 1994, and rapidly gained use for technical documentation which had been notably absent until then. This evolved with departmental web-servers and an intranet site for the administration and scientific programme management, each with its own support staff.

Final Comments

Initially neutron scattering employed state of the art numerical control to provide 24/7 operations for precise mechanical diffractometers and spectrometers to maximise use of the feeble intensities available. By the time the ILL was created paper tape and electromechanical comparators evolved into small control computer systems individually crafted to perform the same functions. Time of flight techniques developed from event recording on magnetic tape, through multi-channel analysers, to using small control computers.

A tour de force was achieved in 1975 at the ILL in measuring the inelastic scattering from ^3He on IN5, which required more than ten days counting without re-zeroing the counts. A few years later the spectrum could be seen after an hour on IN6, which still only had detectors covering 3% of the solid angle.

As detector systems evolved (q.v. the second souffle) data volumes increased. Neutron studies still used important features of the absolute nature of the intensity measurements to perform fundamental experiments. Each year the Vercors group tried to fund an additional 1 square metre of detectors for IN5.

Spare computing capacity and general availability of Fortran on control computers lead to scientists gaining the ability to exam data immediately, and perform better experiments, especially where it was necessary to extract the science from multiple measurements. In diffraction the completed sequences of scans were analysed on the central computer.

In the mid 80s the advent of 32-bit computing on instruments offered a wider choice for distributed computing, ending with ethernet connections around the Institut.

Data treatment was initially performed on even more expensive main-frame computers. As graphics evolved the value of hardcopy output from the original storage scope Tektronix terminal was more and more appreciated. Raster graphics became possible as memory prices dropped and microprocessor speeds increased. Colour graphics using domestic microcomputers became a reality when these used colour lookup chips to generate the colour raster; 35mm polaroid slides were used for hardcopy.

The supporting staff for these activities diminished. Machine code and assembler programming disappeared from instrument control, these tasks being performed by libraries which later added necessary networking features in the current era. For the central service, the operators have been relieved of tape loading and physically moving data now that many terabytes of storage are available on-line in a shared database. Many scientific programs have been introduced from external laboratories and universities. Overall the residual programming activity at the ILL needs little mathematical support.

In contrast assistance is required and generously offered to support the very large park of personal computers, including Macintosh, Microsoft and Linux systems, together with the network infrastructure. Fortunately the hardware reliability of most of these single-board systems is very high.

Because of the impact of computers on every aspect of the ILL the total investment in hardware and manpower for computers, networks, and control electronics is still a major item of expenditure. Data treatment is less rigorous; there are greater volumes, and the scientific result can often be shown simply presenting pictorial evidence of effects of external scans of temperature etc, on more or less raw data. Absolute cross-section measurements, e.g. isotope variation for liquid structure work, is rarer, and fewer such experiments are attempted than in the past. Hence the value of the graphical packages included with the commercial software (IGOR, Matlab etc), though a considerable effort is still spent on adapting the GUIs for daily use by visitors.

Today digital control is everywhere, at low cost, using networked microprocessors which can easily handle parking a car, and even driving it safely. The solution to any computer related problem, from mobile phone to personal computer is well known to everyone; a reboot will, hopefully, restore normal operations!

Ron Ghosh, Epsom, November 2017

Appendix I Central Computer Replacement 1979-91

A more heavily political choice was that to replace the DEC-PDP10 by 1980/81. Already over seven years old it was heavily used by all groups. Competitive

tendering lead to offers from IBM, ICL, Siemens, CII-Honeywell-Bull, and DEC.

Configuration 1978

PDP1070 processor 36-bit (1973), hardware paging, ~256kwords 36-bit memory
4 RP02/3 40 Mb disks, 1 Ampex compatible 40 Mb disk (removable/interchangeable packs)

4 800 bpi NRZ 9 track tape units

2 DECtapes (7 track 0.75 inch block replaceable tapes 2.5 Mb)

1 paper tape reader (boot device)

DECwriter LA30 console, VT05 operators monitor

2 GT40 graphics, 2 Tektronix 4010 graphics+hardcopy, Calcomp plotter, 2 LP10 printers

~10 user terminals - ASR33 teletypes 10cps

3 Network instruments (D11, D17, IN6) ILL protocol; serial line 9600baud)

The performance of the KI-1070 machine cpu was similar to the 1978 VAX11/780, though the disks were connected through channel hardware. (Later VAXes used proprietary interface, the BI bus interface chip, offering better performance than the basic UNIBUS.)

New machine requirements

Approximately 3-4 times processor speed.

Disk capacity 2 Gb

Network front end processor/s

System cost (excluding man-power for program conversion) 12-16MF (1980)
(6.8-9M€ 2017 value)

IBM

A virtual memory system, 32-bit VM-CMS system was proposed. A major weakness was the necessity to pre-allocate disk space to users, rather than the more common and flexible pooled space of the other systems. Networking to the existing ILL systems required major development to adapt to the IBM proprietary SNA system.. At that time there were not many systems in use elsewhere. With two American companies this proposal was eliminated at an early stage.

ICL

An ICL 2980 32-bit system was proposed by ICL running VME/B system. This was used in several UK universities, with varying degrees of success (there was considerable discontent at Oxford). With a commercial background the system mostly used intelligent data entry terminals (non-graphical) which performed local editing etc, rather than simple ASCII terminals. To satisfy the data disk configuration required (2 GB) required 16 disk units. The new computer room was dimensioned sufficiently to contain such a large system. The sales group had its own technical group, effectively isolating the development engineers from any contact with clients.

Siemens

The Siemens BS2000 32-bit (IBM machine code) system was proposed by this

German company. When a number of practical problems were identified by the ILL, Siemens staff were very willing to introduce changes to simplify introduction and use.

CII Honeywell-BuLL

A Multics system running on a bi-processor DPS2 36-bit computer, similar to that at the university of Grenoble was proposed. The system was completely virtual, mapped between memory and disks, and was almost completely written in the high-level language pl1. There had been problems at the university without a guaranteed power supply, and the inevitable electrical storms in summer. This was resolved by copying systematically modified pages, or segments, to tape. Rebuilding the system by system engineers still required several hours. The floating binary format was distinct from the 36-bit PDP10, without the 128 exponent bias; there were hence two representations for 0. As in Unix systems there was a similar hierarchical file structure, and a command given was interpreted as a valid pl1 instruction to seek out and start execution of a segment so named.

DEC

Initially a bi-processor KL1099 was proposed, with similar software to that already in use at the ILL,, offering a simple transition from the existing system.

The sales groups from each manufacturer visited the ILL repeatedly; ICL used the ILL to train its sales staff with unnecessary visits, and hosted a visit to Luxembourg to the EU datacenter to try out their system (and intelligent terminals) with sample ILL programs. We were given access to a Multics system in the US using the early ARPAnet, using a dial-up 30cps modem. We moved several programs using the Texas instruments silent 700 terminal with cassettes for file transfer.

A set of programs were made available to the manufacturers who were requested to compile and run the programs with a log file of terminal response times for each operation for benchmarking. These files were then examined at the ILL; there were clear difference in system response times for simple environment assignments to the time taken to execute compute tasks within programs. Some long calculations were included to assess innate computing power.

There was considerable pressure from the British to turn to ICL; similarly the French emphasised the complementarity with CICG in Grenoble. To quantify the cost in manpower for changing systems it was decided to examine twenty programs from the regular experimental activities. Each manufacturer was provided with tapes of source, and solutions, which included Tektronix 4010 graphics. They were invited to ask for help at any stage. The transformation exercise was set to last some twelve weeks, at the end of which they were invited to demonstrate their programs at the ILL.

ICL allocated two average programmers, of a quality which might be expected to aid a transition over a period of up to a year. CII introduced a young American, to work on the CICG system. We never met the whole team which Siemens allocated to the project.

At the end of the trial period representatives all met up at the ILL. Results were recorded on a blackboard. Three programs were accepted as being untransferable since they used DECtapes, which would have to be phased-out.

ICL brought a 2400 baud modem; one programmer ran the programs successive at the ILL, while his colleague in the UK managed any changes. The Siemens team (lead by a daunting Frau Rechner!) brought several terminals and 2400 baud modems and ran their tests in Munich, with a spare "hot" computer available in case of problems. Unfortunately the Multics at the University had failed in the morning; there was much reduced time for the young Whizz-kid to show his work.

At 1600 Frau Rechner (Siemens) insisted the room with the blackboard be closed!

In practice the initial results showed all machines were technically acceptable, if program transformation could be funded. The usability of the machines could be judged by comparing the scripts for running programs which were actually proposed in the tests. As an example, Siemens and ICL could not explicitly open files within Fortran programs. Siemens were prepared to add this feature. ICL introduced a pause into the program, where the user could enter the required association of the filename and I/O stream, then continue execution. To facilitate use ICL effectively used command files to hide the system complexities, but scripting these also had to be coded. Over the long period of program tests ICL identified that a considerable fraction of total power was expended in terminal handling for their intelligent terminals. By eliminating this code to treat the simpler ILL alphanumeric and graphics terminals this was reduced by 60%, and enabled them to offer a cheaper ICL2972 model, rather than the 2980.

At the autumn steering committee the Directors actually had available one proposal for CII and one for DEC, but neither was actually presented at the meeting.

DEC then announced a new model, the KL1091S, and new disk sub-system, which would be manufactured in the UK (at Ayr). The newer configuration was priced significantly lower than previously, partly due to the inclusion of two RP20 disks, each with 1GB storage on two fixed spindles.

Finally a decision was made to acquire the updated DEC system.

A minimum configuration (cpu, RP06 600MB disk, 2 tape-drives) was delivered shortly afterwards, and housed initially in the ex-NICOLE computer room while the new computer building was being completed. During this period of several months when both machines were operational the final installation was prepared. The 800bpi tape library was rewritten using the PDP11/55 concentrator which had two dual density 800/1600bpi tapes.

Final configuration

DEC 1091S 1M word memory, 3.4 x speed of KI 1070,
PDP11/40 front-end 2RX01 8" floppy disks (boot)
2 RP20 winchester disks (2x1GB), 2 RP06 disks (2x180MB)
4 self-loading tape units (1600/6250 PE tapes)

2 DN20 communication processors (actually PDP11/34)
2 LP20 1200 lpm printers
Benson Graph pen-plotter
LA120 printing console
Slow peripherals from the KI10 were reused, including 2 LP10 printers.

Final cost about 8MF (1980) (~4.5 M€ 2017)

Approximately 16 VT100/PT100 alphanumeric terminals were acquired for users
4 Pericom GRAFPack raster graphics terminals

(The MICOM/SATELCOM terminal switching system was installed and operational by the end; this started the tendency to cable up the scientists offices for terminal access, and losing the sharing aspects imposed by the two rooms of terminals.

The computer building

The first PDP10 had been housed in essentially an adapted corridor. A new building was envisaged for the second souffle for EMBL and ILL biology and chemistry, and for computing. Initial designs for computing lacked an entrance towards the guide-hall and ILL4. Because of the laboratories in the Biology wing, the windows in the computing building office section were all above normal eyeline. In addition, to cut costs, apart from the machine room, the walls were poured concrete, and thus difficult to adapt subsequently. A magnetic tape library was placed adjacent to the machine room, and a large storage capacity in the basement allowed for buying paper once per year at a sizeable discount for a complete lorry load.

The size of the machine room and infra-structure had been set by the possible requirement to house the ICL2980 computer. This required about 30m² for the 16 disk packs alone. An uninterruptible power supply of 95KVA for up to one hour (though only about 30 minutes were ever achieved) used accumulators, and a zero-current switching solid state converter was located with air-conditioning for a similar power load. A Halon gas-extinguisher system was installed in the single open machine room. The machine room was designed with a northerly aspect to minimise summer heating. Half the offices faced due south, and were very hot in summer.

An accidental triggering of the Halon system blew the doors of the machine room open. Revised door mechanisms were put in place, but about a year passed before the Halon system was replenished, due to the costs involved. A source of vibration throughout the building at 7Hz was finally located in the air conditioning system where pipes from the compressor had been cemented into walls, transmitting these low frequencies.

In later years when the power and space requirements reduced the operators were first separated from the noise and cold of the machine room by a partition. Later the uninterruptible power was shared with other technical areas in the institute, especially once unix systems were introduced. finally the air conditioning was shared with the remainder of the building, though this involved creating ducting

through the poured-concrete walls.

Appendix II The conversion to Unix

The proliferation of DEC computer systems had always been highly questioned, primarily by the French. To break this monopoly a change to Unix was imposed in 1992 during the shutdown of the reactor. Very limited financial resources were available for practical implementation. An extensive study within the Département Informatique, with input from external advisors led to a document, The Unix Plan, to cost and implement the transformation.

Two computer systems, the HP 700 and the SGI computers were chosen. The latter, who finally owned the MIPS processor chip manufacturer, had targetted the scientific community, taking a significant part of the DEC market for scientific as the VMS systems showed shortcomings in graphics and overall computing power. At this stage the crystallography community was developing rapidly using the advanced SGI graphics (OpenGL etc.). The HP computers were powerful, but the software support was not driven strongly in the scientific domain.

Up to this time LAN networking used DECnet, and the DEC products DFS for filesharing andDQS for sharing print services. WAN was restricted to X-25 networks, which could also carry DECnet.

With Unix there was a shift to adding IP services to the VMS systems, especially the central VMS servers. The initial use of Ultrix on a DECstation with a MIPS processor was not extended to other systems.

By 1994 each scientific group had at least one workstation available. The Three-axis group, time-of flight and Nuclear Physics were given HP-730 systems. The LSS group and Crystallographers were given SGI-Indigo workstations. An SGI-Challenge-L twin cpu system was ordered for a central system and data server, following examination of systems at CERN. At this stage the VMS service was performed by a cluster of an AXP3000 and AXP3500; these had quite extensive SCSI DECstorage farms for disk storage. These were later integrated partially into the newer systems.

The database for experimental results depended on DEC-Datatrieve on the VMS systems, and most data was stored in VAX11 binary. A decision was taken early to refer to a simpler filesystem approach to store these data, and to use ASCII files to avoid problems of binary/little-endian/big-endian compatibility. The basic format for ASCII data export since 1982 was the fixed length records of the TAPDAT format, which was easily extensible and was adopted, with a few minor additions for the unix systems. To save space, and offer better subsequent network utilisation the ASCII data were compressed using an open source variant of the unix compress tool on the VMS systems. There was an advantage on the VMS systems for using fixed length records; this allowed big files of time of flight data etc, to be read as direct access on VMS. For the unix systems the huge file caches built in to the unix systems made this less necessary, especially given the faster processors. The fixed length format was always easy to read at a human level. There was an opportunity to add commentary zones on the contents of the metadata; this was done for many

of instruments on the conversion of the instrument computers to unix.

On restarting the reactor there were fewer instruments due to a reorganisation. Most still were using DEC computers, and the data conversion was performed on the central VMS cluster.

The staff departures, encouraged during the long shutdown, resulted in a number of scientists busy in software development for their own use being markedly diminished as they either retired early, or left for active science elsewhere. New replacements (in 1995) were not used to the “crude” third generation software tools and adopted packages, notably IDL, MathLab and Igor, which provided GUI interfaces and basic mathematical functionality. These proprietary systems were expensive, and initially the programs could not be exported easily, unlike the halcyon days of the VMS systems, where the executables could be shared and run anywhere. A great deal of time was spent on the GUI, rather than developing the underlying scientific analysis.

The rapid drop in price of workstations, especially after about 2000, when Linux started to be introduced on desktop PCs lead to the unheard of power of offering each scientist his own workstation. From a total computing capacity of some 500 MIPS on the Central AXP cluster in 1993 it was possible to envisage each scientist having similar computing power on their own desks. Some software systems were examined, where this unused power at night could be shared, but for the most part the type of computation done at the ILL were not so time consuming.

Crystallographic problems which had run overnight in the 1970s could be completed within a few minutes on modern machines.

Typical RISC CPU speeds in 1995 were 300-500MHz. It is more difficult to compare with the VAX CISC instructions which rarely exceeded 20MHz! Graphical interfaces were capable of huge consumption of this resource; two X-window scrolling terminal windows which overlapped required considerable cpu power to evaluate whether it was necessary to update every pixel in the two windows. With the advances of graphics processors under Microsoft Windows there were even more attributes per pixel to consider, and again much of the processor was consumed in this activity, as was most programming effort. After the restart in 1995 instruments were changed to unix using DEC OSF AXP-114 computers, then twin-cpu IBM-PCs running linux; electronics were addressed through ethernet links.

Appendix III Approximate computer and component costs

Initially the high cost of computers inevitably lead to the requirement for a number of highly competent staff to maximise their usage.

The custom-built Nicole and Carine systems must have cost a lot at the time, far more than the standard central DEC system, which was finally acquired after 1975, when third party peripherals could also be added.

Very approximate purchase prices (from memory!) are shown with post-inflation conversion to 2017 Euros in brackets.

1979 Tektronix 4010 graphics storage screen terminal 70000F (36500€, 2017)

1981 The DEC-1091S replacement cost about 8MF: 1Mword, 2GB fixed disks, 4 tapes etc. (3.3M€)

1981 VT100 alpha-numeric vdu terminal 10000F (4100€)

1984 VAX11/730, 2MB memory, R80 128MB fixed disk, RL02 10MB Cartridge disk 2 DECTape-II cassettes (for 8085 boot processor), LA100 console
The cost of the first machine bought for initial software development was 358358F (1984) (107k€ 2017). With the potential of three further acquisitions (D11, D17, D22) DEC reduced the price for D11; a second-hand 11/730 was found for D17, and D22 was finally fitted with microVAX-II.

1985 1MB ECC 32bit VAX memory 10000F (3000€)

1984 the PDP11/24 system for IN20 with RL02 disks, graphics terminal etc cost about 400kF, (119k€)

1985 IBM PC-AT cpu 50000F (13900€) (640kB, 20MB disk 2x5" floppies), Hercules Graphics (720x350 pixels) 15000F (4100€), Keyboard 10000F (1300€)

1985 Macintosh plus 50000F (13900 €)

1986 ATARI 1040ST 4000F (1150€) (640x200x16 colours, 1MB memory, 1 3.5" floppy), Sony 13" trinitron monitor 2500F (700€)

1987 ATARI memory 1MB 1600F (500€)

1987 IEC11 GPIB (IEEE-488) UNIBUS interface 33000F (8400€)

1988 LN03R PostScript printer 72000F (17800€)

1988 DEUNA Ethernet UNIBUS interface 30000F (4500€)

1988 microVAX 3500 (IN5) 400000F (99000€)

1989 vt220 alpha-numeric vdu terminal 2500F (600€) (these were bought in packs of 5!)

1989 VAXstation 3100/30 8 MB, 128MB disk TK50 tape ethernet 14" display 100000F (24000€) (D11,D17..)

1990 (D22) s/h microvax-II 100000F (24000€), VAXstation 90000F (21000€)

1993 SGI indy, 20" CRT 140000 (29900€)

1993 Mac powerbook 180 33MHz 68030 4MB/80MB, 9.8" greyscale TFT 3870\$

1999 SGI-O2 64 bit workstation 110000F (21300€) (no monitor)

2000 Sony VAIO PGC-Z600TEK laptop 600MHz/256MB/20MB disk 3000€ (3800€)

2002 IBM XEON workstation 3GHz cpu/4GB memory/200MB disk 4000€ (4939€)

2002 Samsung 19" TFT display 2500€ (3000€)

2005 Sony VAIO laptop S4-XP PC 2GHz/1GB/100GB disk £2000 (2324€)

2008 MacBook white 13" 2.4GHz Core2/1GB/ memory120GB disk £800 (1250€)