EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH

CERN-Data Handling Division
DD-79-10
8 October 1979
E.M. Rimmer

REQUIREMENTS FOR STANDARD PRACTICES
IN FUTURE LARGE HIGH-ENERGY PHYSICS EXPERIMENTS

To be presented at the
Nuclear Science Symposium
San Francisco, 17-19 October 1979

REQUIREMENTS FOR STANDARD PRACTICES IN FUTURE LARGE HIGH ENERGY PHYSICS EXPERIMENTS

E. Margaret Rimmer
Data Handling Division
European Organization for Nuclear Research
1211 Geneva 23, Switzerland

Summary

During the past 18 months, over 500 physicists and engineers have participated in a Working Group, set up by the European Committee for Future Accelerators (ECFA), to study a project for a Large Electron Positron colliding-beam machine (LEP), in Europe in the 1980s. 20 specialized subgroups have investigated technical questions in various areas, and the likely impact of LEP on subnuclear physics and the physics community.

The data-handling subgroups have concluded that techniques currently available, with continued development, will be adequate to meet the technical requirements of LEP experiments. However, there is an urgent need for standard practices at all levels of data handling, to ensure that experiments can be carried out by collaborations involving large numbers of physicists from many institutes.

Introduction

In May 1977, ECFA recommended that LEP studies should be carried out for a machine to be built in Europe in the 1980s. In 1978 a Summer Study was held at Les Houches and CERN¹ and a Design Study was published for a 15 to 100 GeV machine². An ECFA-LEP Working Group of about 300 physicists and engineers, mainly European, was set up to consider the over-all technical feasibility of the proposed machine and its impact on subnuclear physics and the physics community. During 1979, a second Design Study has been published with an up-graded beam energy of 22 to 130 GeV³ and the Working Group has prepared a report for ECFA⁴. In parallel, another ECFA Group has been studying sociological aspects of High Energy Physics (HEP) activities in the CERN Member States and will report early in 1980⁵.

The ECFA-LEP Working Group was divided into 20 Specialized Study Groups and this paper presents, in part, the findings of those concerned with Data Handling⁶. Whatever the final recommendations of ECFA will be, our work has clearly focused an already evident need, namely that standardization is required in many areas to enable huge experiments to be carried out by collaborations involving hundreds of people from tens of institutes.

We have recommended to ECFA that a small group should study further the kind of standards needed and how the HEP community can be encouraged to adopt them. This initiative comes from the physicists themselves, who are already familiar with the problems of organizing large experiments. Therefore, we are sure that the group will receive encouragement and that its recommendations will be supported.

Data Loads

The latest LEP design is for a machine of 30 km circumference, with 8 equidistant beam intersecting regions located around it. The predicted luminosity is $10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ at 86 GeV. There will be 4 circulating bunches per beam, giving a bunch crossing frequency, and therefore a primary strobe rate, of around 50 kHz (20 ns) at each intersection.

Interesting physics event rates are calculated to vary between a few per day for the production of charged intermediate vector boson pairs W^{+-} , to 1 or 2 per second at the peak for the singly-produced neutral Z^0 .

These events must be extracted from a background of unwanted events by a sequence of selection and filtering processes. Average multiplicity per event is expected to be 25 clustered tracks, recorded in a high-resolution detector. Particle identification and momentum resolution will require about 150 samples per track, with typically four 16-bit words per sample for wire address, drift-time, charge ratio and dE/dx data. There will be data from ancillary detectors, such as calorimeters and scintillator hodoscopes, and the estimated total data volume is around 40 kbytes/event⁷.

These data rates and volumes are similar to those in experiments already running, for example at PETRA, or soon to begin, for example at the CERN $p\bar{p}$ collider. They do not, in themselves, present any insurmountable technical problems. Techniques for selective triggering, front-end processing, data compression and correction, event filtering, mass storage and so on, are either already adequate or on the way to becoming so.

However, the high volume and complexity of the data means that the capture, storage, and processing must be handled more "professionally" than is often the case at present. A set of standards is needed early in the design and planning stage, at least within the confines of any one experiment to ensure internal consistency, and preferably with wider validity.

Socio-Economic Factors

The size and cost of LEP experiments will result in a large fraction of the European HEP budget being devoted to experiments at this machine. High-Energy Physics activity will centre around a very few huge experiments in preparation, production, or analysis. It will not be possible for physics departments to have a large fraction of their staff absent for long periods, nor can they give up the design, research, and development work usually done at the home institute. Also some work loads, for example off-line analysis, are so large that part will have to be done elsewhere than at the LEP laboratory.

Methods must be found for a large fraction of the work to be distributed amongst the collaborators and carried out at home. This is especially critical for small institutes, which otherwise will find it difficult to keep up with HEP. These institutes train many physicists, and interest in the subject will decrease if it is impossible to participate. The same argument applies to the smaller Member States of the CERN organization, whose HEP budgets may be quite modest. The work, for example building hardware or writing analysis programs, can only be distributed within a well-defined framework, so that the separate parts will be plug-compatible. In addition, good international data-transmission network facilities will be needed to handle the enormous communications requirements.

LEP collaborations will undoubtedly attract physicists from continents other than Europe. Therefore, it is essential that proposals for standardization should have a wide HEP audience, so that many opinions are taken into account.

Standardization at CERN

To demonstrate that standards can be introduced with success into HEP, we give a brief summary of the

situation at CERN, where experience has been very positive. We also indicate present activity aimed at satisfying future needs.

The international structure of CERN means that standards specified internally rapidly spread through the European physics laboratories and elsewhere via visitors, and adoption of external standards enables collaboration between physicists from the Member States. It goes almost without saying that this saves an enormous effort in design, documentation, and maintenance. A CERN survey has recently been compiled covering the use of ISO and IEC standards⁸. The publication was accompanied by an appeal from the CERN management for the active promotion of these standards at CERN, demonstrating that a positive attitude is being encouraged throughout the laboratory.

The physics, electronics and computer groups at CERN have adopted many standards defined by bodies such as NIM, ESONE and IEEE and have often participated in committee work to define them, for example in the ESONE CAMAC groups; this is continuing into the next generation of hardware and software specifications, for example in FASTBUS studies.

For data acquisition, CERN has had tremendous success with standardization. With the advent of CAMAC, software standards became feasible and a project was launched to provide facilities for the Hewlett-Packard 2100 computers⁹. Two CAMAC packages were written; one for executing individual CAMAC operations (CAMAC calls) and one for interrupt-driven data acquisition¹⁰. All subroutines are identically callable from FORTRAN or BAMBI, a CERN dialect of BASIC¹¹. The CAMAC calls are intended primarily for hardware testing and debugging and are usually run interactively in BAMBI. The data acquisition routines are intended primarily for fast data taking, buffering and storage, and are usually run in FORTRAN. A BAMBI/FORTRAN translator is available, so that users can create their program interactively in BAMBI for ease of debugging and convert to space/speed efficient FORTRAN for the final version.

This HP-CAMAC facility was developed over a period of 5 or 6 years and is in use on over 50 systems at CERN and elsewhere at any given time. The decision to produce modular software, which is configured with the physicists for their particular application, has increased support personnel efficiency from around 1 system per person, to about 7 systems per person. Programmers can devote spare effort to improvements and development projects and an experiment is not critically dependent on one particular expert. The software is thoroughly debugged, having been used in hundreds of different environments from single-crate test systems to multi-branch experiments.

This type of on-line computing support was very well received by most physicists at CERN, and similar facilities are now available for the Norsk Data¹² and Digital Equipment Corporation¹³ computers which form part of the CERN On-Line Computing Pool. The CAMAC calls are also available in a BAMBI system running on CAVIAR, an autonomous processor based on a Motorola 6800, which includes a EUR 4600 CAMAC driver and an IEEE-488 bus controller¹⁴. The original calls have recently been updated on all machines to conform with the NIM/ESONE specifications for CAMAC subroutines¹⁵.

We now consider some aspects of HEP data handling in areas where further studies should delimit common practices. We firstly look mainly at hardware features and treat the software in detail later.

Hardware and Data

We assume that CAMAC 4100¹⁶ and 4600¹⁷ systems will form part of LEP experiments. Certainly CAMAC will remain adequate for many data-acquisition and storage functions of the on-line computer. The microsecond CAMAC bus cycle matches that of most minis quite well and is sufficient to keep pace with modern mass storage devices, such as 6250 bpi/125 ips magnetic tape. Present work to define 4100 extensions allowing handshaked cycles and extended addressing will add to the applicability of CAMAC¹⁸.

In a HEP data-handling sequence, the aim is to reduce the event (trigger) rate from the bunch crossing rate to something approaching the interesting physics rate¹⁹. The selection of wanted events from the candidates is made in a series of increasingly sophisticated, and therefore increasingly time-consuming, processes. Event data are passed through an inter-connected processing chain from the detector to the off-line analysis program. The format or representation of the data at each stage, whether electrical signals or recorded bit patterns, cannot be independently chosen if efficiency and compatibility are to be optimized for succeeding processes. Global consideration must be given to requirements at all stages in the sequence.

Detector Design

The complexity of LEP events is reflected in the pattern recognition and track-fitting programs. Extremely sophisticated software will be required for these geometrical reconstructions, and in some cases complex interactive graphic techniques will be needed. The more suitable the raw data is to algorithmic manipulation, the more easy it will be to produce effective software. The ultimate goal is to perform full event reconstruction on-line for all events which pass the early selection filters. Even if this goal cannot be reached, a rigorous recognition test will be required to keep the storage rate down to some reasonable level.

Optimized data representation is already carried out in detector design, such as the choice of angles for wire-chamber planes to produce integer trigonometry. It is clear that not only improved pattern recognition techniques are needed, but also a closer collaboration between detector designers and geometry programmers, to produce more detector features which match the parametric representations of the software²⁰.

Primary Triggering

With the increased use of fast Emitter Coupled Logic (ECL) integrated circuits in NIM and CAMAC modules, there is a move towards ECL front-panel signal levels. Also, to reduce cost and cable volume for large experiments, twisted pairs are replacing conventional coaxial cables. Emitter Coupled Logic technology should cover LEP requirements for fast signal transport, primary trigger decision logic operating within about 100 ns after the bunch crossing strobe, and front-end data processing. The logic should reduce the event rate by more than one order of magnitude.

In the absence of ECL standards, there is already a risk of incompatibilities. A CERN group has been set up to find a compromise between the different ECL systems being used in the laboratory²¹. It is hoped that their recommendations can form the basis of a more widely accepted standard for cable/connector conventions, especially for new systems such as ECL implementations of FASTBUS cable segments.

The group will next define a basic set of modules, such as coincidence units, to be designed for the ECL environment. NIM-style modular logic in ECL, particularly with multi-line signals, should be well suited for HEP primary and secondary triggering²².

It is also necessary to minimize the variety of detector read-out systems. They should be designed for general application whenever possible, as is, for example, the CERN ECL Receiver-Memory-Hybrid module system RMH²³.

Secondary Triggering

Secondary triggering consists of combinatorial logic of greater complexity than in the primary trigger. It requires more input data and takes a few microseconds to operate. At LEP, this decision can take up to 20 μ s without reducing effective luminosity. It should therefore be possible to reduce the event rate by at least another order of magnitude, leaving only a few tens of events per second to be handled subsequently.

At this stage in the data-handling chain, semiconductor memories and microprocessors begin to play a role. For example, correlation tables can be dynamically loaded into read/write memories and with fast table look-up techniques can produce powerful and flexible secondary triggers. These and similar methods are already employed in several experiments (Ref. 21 gives some CERN examples). These devices should be standardized to the extent of being housed in the FASTBUS system, which allows access from a host computer and point-to-point communication anywhere within the system itself²⁴. The logic should be modular, so that several modules can be coupled together to increase the efficiency of the decision process.

Trigger Processing

Trigger processing is the next filter applied to the data, during 100 μ s or so after the secondary trigger, depending on dead-time constraints. It involves a primitive track fitting or similar pattern-recognition procedure, and is optimally implemented in a fast stored-program device for flexibility. These devices should also be housed in FASTBUS. Not only does the bus speed match that of bit-slice processors, in the 100 ns range, but a well-integrated triggering sequence can be more easily created if the associated processors are intimately connected. We need FASTBUS systems as soon as possible, to avoid one-off systems at the front-end of experiments; otherwise experience will not be cumulative and incompatibilities will eventually become overwhelming.

Another result of adopting standards for various logic functions, module characteristics, and so on, is that if an item is reproduced in large quantities, industry might manufacture it as a pre-prepared kit, integrated circuit, or other package. For CAMAC, boards can be purchased with pre-wired Dataway driver logic; floating-point operations are now almost nowhere performed by software. If functions and features of HEP data handling can become commercial components, there will be an increase in reliability and decrease in cost.

Event Filtering and Equipment Monitoring

Programming power and flexibility rather than speed are of prime importance at the next stage. For event filtering and equipment monitoring, highly complex pattern-recognition procedures and maybe even a full physics analysis are needed. For event filtering, one or more powerful minicomputers, and/or microprocessors of equivalent capability, will be necessary. For full equipment monitoring, a connection to a larger remote

computer may also be needed. This latter technique is currently being used at CERN by, amongst others, the European Muon Collaboration, who have a PDP-11 cluster connected to the central IBM 370/168 via CERNET²⁵.

Microprocessors built to emulate the operation of well-known computers, such as the 168/E²⁶ and MICE, which emulates a PDP-11²⁷, will be increasingly used to run powerful on-line filtering programs. Emulators with the facility to combine emulated code for ease of programming with microcode for efficiency of execution, for example MICE, will be ideally suited to this application.

Emulators can also be used off-line for running more complicated programs²⁸. Inefficiencies are compensated by low cost, and events can be treated individually and identically by a set of emulators. This simple parallelism involves no complicated inter-task, inter-processor communications; speed is gained almost linearly by adding more processors. To fit well into the several bus systems used in HEP experiments, these emulators should be equipped with FASTBUS and/or CAMAC ports. Such ports are already foreseen in the two emulators mentioned above.

Data Acquisition

Towards the end of the data-handling chain, the event data should be collected in a buffer, suitably formatted for mass storage as records in a file. The FASTBUS system is ideal for such data collection from scattered source registers.

The connection of FASTBUS to a host computer can either be via a FASTBUS/host interface, or indirectly through a FASTBUS/CAMAC interface. This host will organize data collection from both FASTBUS and CAMAC, local data storage, and link connections to other processors. The arrangement of the stored data within the files will depend on the needs of the off-line analysis.

Networking

The logically independent on-line tasks of a large HEP experiment are often implemented in physically separate computers. This requires links for multi-computer clusters at the experiment and between the experiment and a remote centre. ECFA-LEP studies of future HEP scenarios showed that this trend will accelerate.

We have had typically painful experience at CERN, where many links -- serial, parallel, CAMAC, non-CAMAC -- have burgeoned over the years, with the accompanying difficulties of support and maintenance. Links are particularly tricky to service, as they often join dissimilar systems, rarely both well-known to one person. To improve the situation, we have started a project to standardize all links to experimental computers using CERNET hardware and protocols. The hardware is a double-width CAMAC module with a 1K word buffer of 16-bit words²⁹; line transmission is serial, at up to 5 Mbits per second³⁰. The system was built before suitable international standards had emerged and uses CERN protocols. The STELLA satellite link between CERN and other laboratories will use the HDLC line protocol³¹.

We are considering software written with ESONE/NIM CAMAC calls embedded in a high-level language such as FORTRAN. Another attractive possibility is to implement the software in a programmable Auxiliary CAMAC Crate Controller, developed at CERN and based on the TMS 9900 microprocessor³². The user program interface to this Controller could then be at a high level, and a large fraction of the space and time presently required in the computer for link software could be eliminated.

Other facilities could be standardized as firmware in CAMAC controllers if wider use is made of the CAMAC 6500 standard³³, which supports multi-controllers on the Dataway. Greater use should also be made of the CAMAC serial system³⁴ for long-distance, slow data traffic, such as infrequent control and monitoring functions. This would rationalize the integration of these subsystems into the over-all data-acquisition system. Serial and multi-controller CAMAC have been largely ignored in HEP. They are widely used in accelerator control and must be adopted by physicists as experiments begin to have similar control and monitoring requirements.

Software and Languages

We now look more closely at the software environment. People can have real difficulty adjusting to more than one type of computer system or language with a bewildering variety of rules and conventions. Efficiency rapidly diminishes if there are different ways to perform the same job, such as character rub-out on a keyboard. Almost superhuman organization is required to maintain different implementations of what is functionally the same program.

LEP physicists will be confronted with a whole collection of processors simultaneously, on-shift, where an array of terminals will connect them to various micros, minis, and maxis, in addition to the collection of off-line computers used for equipment testing, data analysis, and other distributed jobs. There will be system and application software for all of these devices, and it will be essential that several members of the team have a clear over-all understanding of the subsystems and their interconnections.

Obviously an interface is needed between the user and the equipment, so that differences between devices or subsystems are hidden or made transparent. The more subsystems which can be encompassed by a uniform interface, and the further such interfacing can be pushed towards the actual hardware, the less complicated will be the total system.

High-Level Languages

Physicists are rarely excited by the attractions of the latest computing language. It is more important that programs can be easily produced, debugged and maintained, and, if necessary, transferred between machines. The languages in which they are written must be widely used by programmers and manufacturers. It is preferable that the language lends itself to good coding practices, such as structured programming. In any case, a set of guidelines should be established, so that software produced independently by members of a collaboration will be jointly understandable, runnable, and maintainable³⁵.

For minis, midis and maxis, two kinds of high-level language are required; an interactive language suitable for equipment testing, and a compilable language for fast data manipulation. The languages used most widely are BASIC and FORTRAN, respectively and there is no change imminent³⁶. It is therefore desirable that an International Standard for BASIC is defined as soon as possible, and that FORTRAN77 is widely adopted to conserve the tremendous investment already made in FORTRAN programs.

Other Portable Languages

Standardization at the intermediate and lower language levels is more difficult because details of processor architecture become increasingly evident. It is usually systems and applications programmers who use these languages rather than physicists themselves.

Nonetheless, with mixed-mini systems, and mini-micro combinations, program production and maintenance are expensive and difficult if portable languages are not used. Some experience has been gained at CERN using the language BCPL³⁷. BCPL is portable by virtue of its compiler, which is itself written in BCPL. Its biggest success has been in CERNET software for the on-line minis and CERNET Modcomps. The software is now being translated into machine-dependent languages to improve performance, but BCPL software will remain the reference version for future updates to eliminate risk of divergence.

Another potentially portable language gaining popularity is PASCAL. The PASCAL compiler is available in PASCAL and has been implemented at CERN as a cross-compiler running on the IBM 370/168 for the TMS 9900 microprocessor³⁸. PASCAL compilers for minis are already produced by the major manufacturers.

Use of Hardware Emulators

Implementing portable compilers or cross-compilers is adequate if the number of host/target combinations stays small. Experience to date gives little hope that this will be the case for micros. At CERN, official support has been limited to Intel 8080s, Motorola 6800s, Texas Instruments TMS9900s and Ferranti's F 100-L.

With the increased use of bit-slices, a language explosion can be avoided if they are configured to emulate well-known computers (see above). Less programming and support effort is needed than for a micro with special machine code, and programs can be produced, tested and debugged in the more friendly environment of the emulated host system. Standardization should automatically emerge if widely-used machines are emulated.

Object Module Format

Another technique to reduce the support effort for lower-level languages is to use a machine-independent format for the output from compilers and assemblers. There is a lot of activity in this area and an IEEE project is devoted to it³⁹. A CERN group have had considerable success in introducing CUFOM, the CERN Universal Format for Object Modules⁴⁰. Compilers and assemblers are available for several low and intermediate level languages, which generate relocatable output modules in CUFOM. These modules can then be manipulated by a universal linkage editor if necessary, to produce yet more object modules. A relatively simple target-machine-dependent loader is then required to produce executable code from CUFOM modules.

Mnemonics

It is too late to hope that manufacturers will adopt a unique set of mnemonics to describe machine instructions, and we must live with senseless irritations such as JSB, JSR and JMS for subroutine access. However, engineers designing programmable devices for HEP should agree on a common set, preferably closely resembling that of a popular mini. Some thought has been put into this already⁴¹ but more work is needed.

There is growing interest in general-purpose processors implemented as list interpreters. This approach is being studied for FASTBUS⁴² to define a bus transaction by loading a FASTBUS module table with a list of words containing address information, data, function codes and so on. The list format and a set of mnemonics should be defined, independent from the processor used in a module.

Stored Data Structures

To facilitate the general handling of recorded data from a large experiment, the file and record format of

the stored data must be standardized so that read-and-unpack software can be written in some standard language and used at all installations.

At CERN we have defined the file structure for unlabelled³ and labelled⁴ magnetic tapes recorded by all on-line computers. The label standard is a subset of the ECMA standard⁵ and is compatible with both the CDC and IBM (ANSI) systems at the CERN Computer Centre and elsewhere.

We are now working on a common format for records and event data within the files. This will be an expansion of the format already supported on the HP and ND computers and at the Computer Centre. It will eventually be implemented also on the PDP-11s, and on the new VAXs. This format must be clearly defined, because it will have an impact on the event data collector and on analysis programs. Features of the analysis software must be considered when defining the format as the data complexity increases, and the introduction of schemes similar to the HYDRA system for bubble-chamber data analysis should be considered⁶.

Libraries

As part of a general rationalization, the CERN on-line group is standardizing calling sequences for all subroutines supported on the minis, compatible with the CERN Computer Library wherever possible. We are starting with histogramming and graphics packages. For graphics, we use various subsets of the manufacturers' software, in this case PLOT-10, and we are presently following developments on IGL, Tektronix's Interactive Graphics Language.

On-line programming is steadily assuming proportions normally associated with off-line work and computer centres. Therefore it is worth while to study some of the management techniques developed for software in these areas. One example is the portable program PATCHY⁷, which allows archiving of program sources, updates and so on, in a systematic way. A record can be kept of program versions and modifications. This is useful both for different implementations of the same program and for chronological developments of a single program.

Communications

We lastly consider communication between the user and the experiment, and between the users themselves. As mentioned earlier, it will be impractical for users to have a variety of conventions for communicating with the experiment. Work is in progress on a CERN standard interface for the simple control functions of experiments such as START, STOP, SELECT EVENTS, and so on. Several international bodies are working on the more difficult problem of standardizing Job Control Languages for operating systems⁸. Some common set of meaningful error messages is also needed, and these messages should be identical on whichever devices they are printed or displayed.

Experience gained during ECFA-LEP work, FASTBUS work and while preparing this paper (on the CERN IBM WYLBUR facility) has clearly shown that members of large, dispersed teams need to be in easy communication with each other. For documentation, messages, etc., it is almost essential to have a central mass storage filing facility, accessible from all participating institutes. With the aid of a good text-manipulator and post-box service, documents can be produced, inspected, edited, and messages can be circulated quickly and easily. In this direction, the FASTBUS group have recently established central WYLBUR files at SLAC.

This work style must be considered seriously by physicists, with more and better use of networks. CERN is investing effort in this field and will soon supplement the on-site CERNET by a gateway to the Post Office networks and a satellite link with several other European laboratories. A leased-line link with the Rutherford Laboratory already provides an excellent service for British visitors and is used both for data transport and for service communications. The availability of these modern services must be extended.

Conclusion and Acknowledgements

Some physicists in Europe feel a strong need for more self-imposed discipline by extending the use of hardware and software standards in HEP. We are sure that this feeling is shared by many colleagues, and that efforts by ECFA members to follow up these initial studies will receive encouragement and support. We have presented in this paper a non-exhaustive, although already long, list of topics which merit our concerted attention, and we strongly urge members of the HEP community to work together to reach some of the goals outlined.

Members of the ECFA-LEP Data Handling group are thanked sincerely for a friendly and fruitful year of work. They are listed by name in the ECFA-LEP Report and their ideas are presented, to the best of the author's ability, in this paper. A special mention is due to Prof. E. Lillestøl, of the University of Bergen, Norway, who initiated the proposal for a detailed study of future HEP Data Handling Standards.

References and Footnotes

1. M. Jacob, Summary report, Proc. LEP Summer Study, CERN 79-01 (1979), vol. 1, p. 1.
2. LEP Study Group, Design study of a 15 to 100 GeV e^+e^- colliding beam machine (LEP), CERN/ISR-LEP/78-17 (1978).
3. LEP Study Group, Design study of a 22 to 130 GeV e^+e^- colliding beam machine (LEP), CERN/ISR-LEP/79-33 (1979).
4. A. Ichichi (Chairman), The ECFA-LEP Working Group Report, to be published, October 1979.
5. J. Mulvey (Chairman), University of Oxford, private communication.
6. Chairmen of ECFA-LEP Data Handling Working Groups: E.M. Rimmer, CERN, Geneva, Switzerland: Data Acquisition. M.G.N. Hine, CERN, Geneva, Switzerland: Data Transmission. R.T. Van de Walle, Nijmegen University, Netherlands: Data Analysis. H.M. Fischer, Bonn University, West Germany: Designing Experiments.
7. F. Niebergall, Note on data loads from LEP experiments, ECFA-LEP Working Group SSG/18/7 (1979).
8. B. Langeset, Use of international standards at CERN, CERN/DG-9/Ed.2 (1979).
9. P. Scharff-Hansen, CERN, private communication.
10. E.M. Rimmer, General purpose CAMAC software for HP 2100 series computers, CERN/DD/78/14 (1978), and Proc. ESO/SRC Conf. on Applications of CAMAC to Astronomy, September 1978 (ESO, Geneva, 1978), p. 93.
11. S. Cittolin, CERN, BAMBI (1977).

12. A. Bogaerts, CERN, The NORD-10 data acquisition system (1978).
13. A. Silverman (Ed.), CERN, The PDP-11 data acquisition system manual (1978).
14. S. Cittolin and B.G. Taylor, CAVIAR: Interactive microcomputer control and monitoring of multi-crate CAMAC systems, Proceedings of the Joint Conf. on Microprocessors in Automation and Communications, University of Kent, September 1978, IERE Conf. Proc. No. 41, p. 309.
15. Subroutines for CAMAC, DOE/EV-0016 (1978).
16. CAMAC - A modular instrumentation system for data handling, EUR 4100 (1972) [TID-25875, IEEE Standard No. 585 (1975)].
17. CAMAC - Organisation of multi-crate systems, EUR 4600 (1972) [TID-25876, IEEE Standard No. 596 (1976)].
18. A. Peatfield, Extended use of the CAMAC dataway, these proceedings.
19. L.O. Hertzberger and C. Verkerk, Notes on trigger logic and processing, ECFA-LEP Working Group SSG/19/2 (1979).
20. H. Grote, CERN, private communication, and H. Eichinger and M. Regler, Preview on a report of track fitting methods, ECFA-LEP Working Group SSG/18/10 (1979). (Full report to be published.)
21. H. Verweij (Ed.), ECL logic for counter experiments, CERN-EP Internal Report 79-06 (1979).
22. F. Bourgeois, Modular trigger logic techniques at the CERN Omega spectrometer, these proceedings.
23. J.B. Lindsay et al., A fast and flexible data acquisition system for multiwire proportional chambers and other detectors, Nucl. Instrum. and Methods 156, 329 (1978).
24. FASTBUS - An emerging laboratory standard, several contributions to these proceedings.
25. D. Botterill, B. Charles and V. White, An on-line data acquisition system for particle physics experiments using PDP-11 computers, CERN-EP/79-52 and DD/79/5 (1979), submitted to Nucl. Instrum. and Methods.
26. P.F. Kunz, The LASS hardware processor, Nucl. Instrum. and Methods 155, 435 (1976).
27. J. Joosten, CERN, private communication, and co-author of Architectural considerations for microprogrammable emulating engine using bit-slices, to be submitted to the 7th Int. Symposium on Computer Architecture, La Baule, France, May 1980 and CERN/DD/79/7 (1979).
28. P.F. Kunz et al., Experience using the 168/E microprocessor for data processing, these proceedings.
29. J. Anthonioz-Blanc and J. Joosten, CERNET network CAMAC link, CERN/DD/78/22 (1978).
30. J. Joosten, the transmission system for the network links, CERN/DD/NPN-31 (1976).
31. The STELLA experiment technical report, European Space Agency (1978).
32. F. Beck et al., Auxiliary CAMAC crate controllers using a 16-bit microprocessor, CERN/SPS/ACC/RR/Report 78-24 (1978).
33. Multiple controllers in a CAMAC crate, EUR 6500 (1978) [DOE/EV-0007, IEEE Standard No. 675].
34. CAMAC serial highway system and serial crate controller type L3, EUR 6100 (1977) [DOE/EV-0006, IEEE Standard No. 595 (1976)].
35. M. Metcalf, Proposal for a set of minimum programming standards, ECFA-LEP Working Group SSG/18/5 (1979).
36. I. Willers (Ed.), High level languages at CERN, CERN/DD/76/2 (1976).
37. S. Mellor and I. Willers, CERN, The BCPL mini-manual (1978).
38. H. von Eicken, private communication.
39. T. Pittman (Chairman), The relocatable object format project, IEEE Task No. 695.
40. J. Montuelle, CUFOM, The CERN Universal Format for Object Modules, CERN/DD/78/21 (1978).
41. J.D. Nicoud, Common assembly language, Microscope 1, No. 7, 11 (1977).
42. M. Johnson, Fermi NAL, private communication.
43. E.M. Rimmer, On-line support for unlabelled magnetic tapes, CERN/DD/OC/79-3 (1979).
44. E.M. Rimmer and J. Ogilvie, On-line support for labelled magnetic tapes, CERN/DD/OC/79-2 (1979).
45. European Computer Manufacturers Association, Magnetic tape labelling and file structure for information interchange, ECMA Standard 13 (1978).
46. R. Bock, E. Pagiola and J. Zoll, CERN, The HYDRA system manual (1979).
47. H. Klein and J. Zoll, The PATCHY reference manual, CERN Program Library (1977).
48. For example, the ANSI Operating Systems Command Language Standardization Committee XSH1, C.T. Schlegel (Chairman), and the IFIP T.C.2 Working Group 2.7, F.R. Hertweck (Chairman).