R. T. Bell
H. Øverås

The Use of Computers at CERN

Abstract: This paper surveys the many ways computers and connected special devices are used in the laboratory of high energy nuclear physics research of CERN (European Organization for Nuclear Research) in Geneva, Switzerland.

Introduction

High energy physics, also called elementary particle physics or recently (and perhaps more appropriately) sub-nuclear physics, is concerned with the nature of and the forces between "elementary" particles, which constitute the smallest known pieces of matter, and of which about a hundred, stable and unstable, have been discovered.

Since only two of these, the proton and the electron, are readily available to the experimental physicist, some special source is needed to produce the other particles. In the early days radioactive elements and cosmic rays permitted some further progress, but the availability of large particle accelerators has changed the situation dramatically, making the research, however, at the same time rather expensive. In order that Europe should not be kept far behind in this research, CERN was created about 15 years ago by 13 Western European states to build and operate first class accelerators for the benefit of physicists from universities all over the continent.

With an accelerator beam a research group may perhaps find as many interesting "events" per second, as it could have done per year without it. Broadly speaking, an "event" is what follows when a fast moving particle impinges on a target nucleon or nucleus. To record all these events, better and faster detectors, such as bubble chambers and spark chambers, have been developed. The large volume of data coming out can be analyzed only with powerful computers, and in addition many of the devices need smaller computers on-line for buffering, logging or control.

A few figures will illustrate the present data handling situation at CERN:

- during the last ten years the total computing capacity has doubled every year;
- an average experiment requires the equivalent of about 500 hours of IBM 7090 time, and many require much more;
- the number of computers at CERN, which started as one in 1958 and was still only two in 1963, is now (visiting groups included) about 30.

Although we try in the following to cover all computer activities (except administration and library) at CERN, the presentation is likely to be biased towards electronics experiments, in which we are more experienced.

In many cases the techniques described have grown out of a more or less formal collaboration with other laboratories, or may even be copied. We make no attempt to distribute credits, but leave this to the indicated references.

General information about CERN can be found in Ref. 1, and about data handling at CERN in Ref. 2.

Detectors

High energy experiments are usually grouped into two classes, depending on whether or not they involve a bubble chamber as detector.

Bubble chambers have, without any change of principle, developed into complex instruments containing many cubic meters of liquid hydrogen or a heavier liquid in a strong magnetic field. The liquid is at the same time target and detector, so a whole chain of decaying particles can be followed in great detail; deflection of the charged particles in the field allows momentum and sign of charge determination. Stereoscopic pictures of the bubbles formed along the tracks of the charged particles are taken with complicated optics. Since this information can be used only after development and analysis of the film, feedback of results during the experimental run is very limited.

The other class, often referred to as electronics experiments, covers a rather wide and rapidly changing range of

The authors are located at CERN, Geneva, Switzerland.

detectors, dominated by scintillation counters and spark chambers of many kinds. In the scintillation counters light produced by traversing particles is "seen" by a photomultiplier and converted into an electrical pulse. These counters are very fast, but give only a rough definition of the particle trajectory, unless each scintillation counter is made very small. They are therefore often used to provide a triggering signal, i.e., to switch on the high voltage, for the slower spark chambers, so only trajectories falling within a restricted space are recorded by them. In the spark chambers the traversing particle causes a local discharge between two electrode planes. The position of the spark can either be reconstructed from photographs, as in optical chambers, or directly digitized, as in the filmless chambers. A type of chamber which is becoming more and more common has at least one of the electrode planes made out of parallel wires (typically 1 mm apart). With some technique using, for example, magnetic cores or magnetostrictive signals, the address of those wires which have participated in a discharge is read out. These wire chambers[3] offer good possibilities for on-line checks and fast feedback to an experiment, and their data taking rate can be high. Besides these detectors, an experimental setup frequently contains analyzing magnets. Usually the spark chambers are outside the magnetic field, but there exist large setups with the chambers inside the field.

Digitization of tracks on film

As mentioned already, all bubble chambers-optical spark chambers produce their output information on film. Bubble chambers normally produce for each event three stereo views on three pictures (frames), and optical spark chambers produce two stereo views on one, or sometimes two, picture(s). The pictures contain, among other things, tracks of interesting events and fiducial marks, and they are subsequently evaluated on special devices, some automated and computer controlled and others hand operated. Figure 1 indicates the devices in use at CERN. Spark chamber events, Fig. 2, are sufficiently simple that the entire scanning and measuring process can be done

Figure 1 Film digitizing devices at CERN.

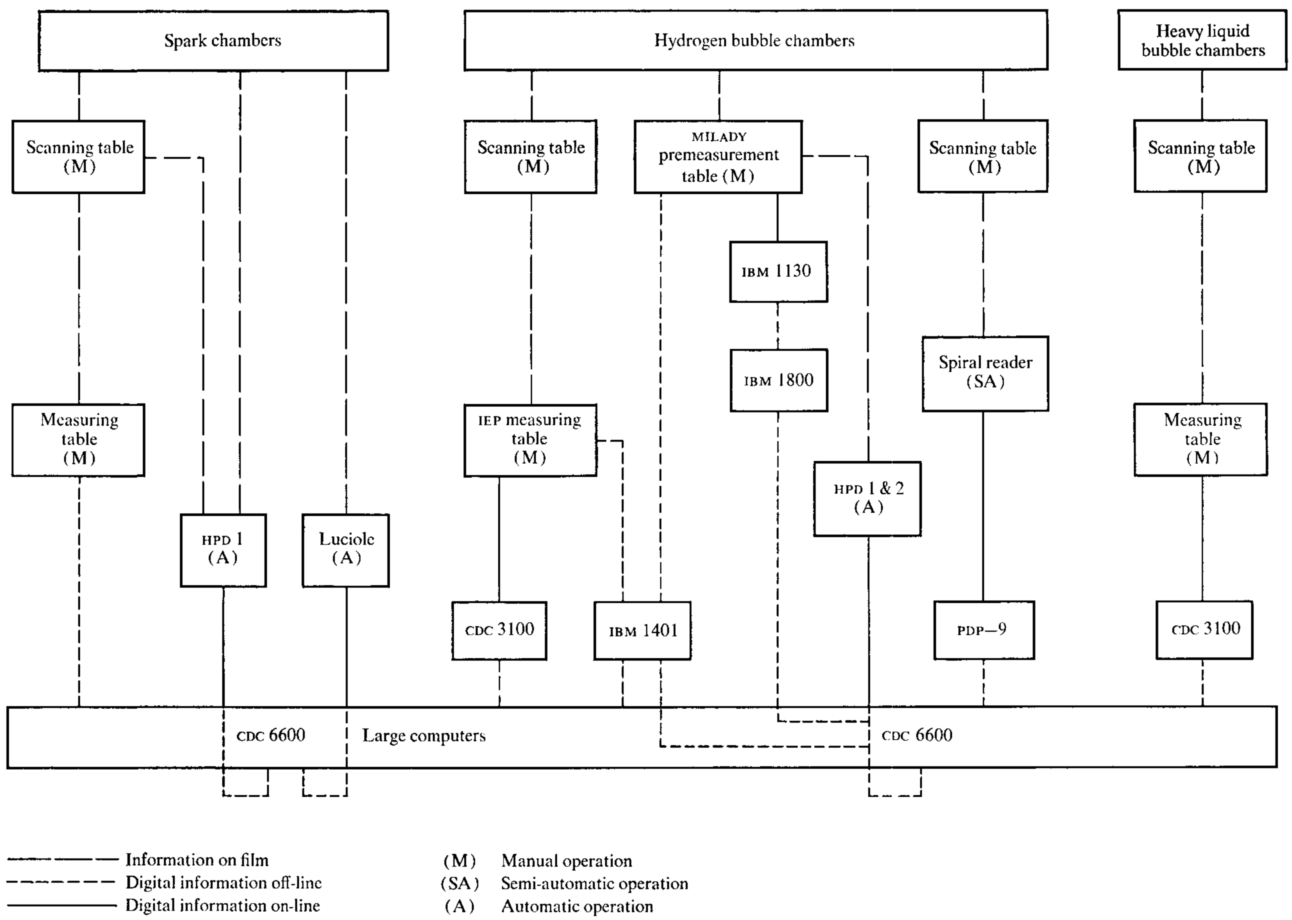

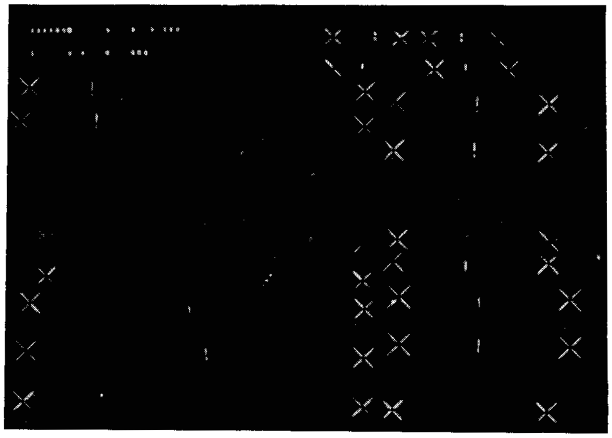

Figure 2 Spark chamber picture. The crosses are fiducial marks, and in the upper left corner is the labelling information. The remaining dots are sparks (or background).

automatically on flying spot digitizers which we will describe later. When, however, the ratio of pictures obviously containing no interesting event is very high, it is worth while to reject these by hand scanning before the automatic processing, since it is hard to improve upon the human eye in economical pattern recognition.

In order to measure the more complex bubble chamber pictures (Fig. 3) on flying spot digitizers, not only scanning but also a premeasurement is necessary to provide guidance information. In fact, a good part of the bubble chamber film at CERN is still both scanned and measured by hand. Both the measuring tables and the premeasurement tables will, however, in the near future all be working on line to small computers, the main advantages being that hardware failures and operator errors can be detected and corrected quickly, so an entire day's production is less easily lost. Use of punched cards or paper tape as intermediate medium is also avoided.

• *Measuring tables on-line*

There are nine measuring tables, called IEP, in use for hydrogen bubble chamber film at CERN. Four of them are now on line to a CDC 3100, and another two will be so by the end of the year. The measured coordinates along the tracks are checked and go on magnetic tape ready for the standard analysis programs. The number of pictures that have to be remeasured because they fail in this analysis has gone down from 15% in the off-line mode to less than 3% in the on-line mode. Six other measuring tables for heavy liquid bubble chamber film are connected to another CDC 3100 with a disk, which permits a slow, but quite sophisticated, geometry program on-line.

• *Premeasurement tables on line*

When the operator, scanning hydrogen bubble chamber film, finds an event which should later be measured on a flying-spot digitizer, he registers on punched cards the number of the picture, the event type, and the approximate middle- and end-point positions of each track of the event on each of the available views. These points serve to define a zone (known as a "road") on the film within which the track will be located. At present one premeasurement table (MILADY) is connected to an IBM 1130, and in a year from now this will be true for five more. The checked data are stored initially on the disk, and at convenient intervals of a day or so copied onto magnetic tape via an IBM 1800, normally used for other work. To avoid this intermediate step the IBM 1130 may later be connected on line to the CDC 6600.

• *Flying-spot digitizers*

The automatic measurement of film has up to now mainly been done with flying-spot digitizers, in which a light spot sweeps systematically over the film frame. Two different types have been made at CERN, one (HPD) with a mechanically controlled light spot, and another (Luciole) using the light spot of a cathode ray tube.

The HPD[4] (or Hough–Powell Device) exists in an old version, HPD1, and an improved version, HPD2. It was proposed jointly by and has been built at several laboratories, including CERN. Optical, mechanical and electronic parts allow the photograph to be scanned by a small spot of light. By splitting the light beam it is possible to scan the film with one spot and a precision picket-fence grating with the other, both at the same time. The signal from a photomultiplier behind the grating allows one to drive a counter whose contents represent the distance of the spot from the start of the scanline. The signal from the photomultiplier behind the film is used to define the time when the spot of light is centered on a bubble or spark image and the contents of the counter at that time give the coordinate of the bubble or spark center. The film is mounted upon a precision measuring stage which is driven at constant speed perpendicularly to scanning spot motion. The movement of the stage is also digitized and its position is read out at the end of each scanline. The scan of a single frame (150×50 mm^2) lasts $\sim$6 sec on HPD2. The average measuring time, including all overheads, is, however, considerably longer.

The main steps in the operation of the device are controlled by means of the program in the on-line computer, the CDC 6600. Instructions from the computer specify when, and by how many frames, the film should be advanced; they also specify at which coordinate of the stage the measuring sequence should begin, at which coordinate it should end, and the separation between scan lines. Thus, once a film has been loaded into the device, it is, in principle, possible for all the interesting pictures to be selected by the computer, scanned by the spot, and the resultant coordinates to be sent back to the computer for each picture, all without the need for any human intervention.

For bubble chamber film, however, control information has to be prepared in advance as described in the previous section, and the program in the CDC 6600 (known as GATE) reads in this information immediately before initiating the measurement by the HPD. During the input of the coordinates from the HPD a first sorting is made. Those lying within the roads, or representing special features such as fiducial marks, are retained and the rest discarded. The points lying inside the roads are then processed off-line in a second program (known as FILTER) which identifies those points belonging to the track of interest and rejects those arising from crossing tracks and background. The output from this program is a set of average points which can be used for input to the same geometrical reconstruction program as used for events measured by hand (on the IEP's). The accuracy is slightly better than for hand measurements, and HPD2 is some 5 to 10 times faster than an IEP.

A system known as Minimum Guidance[5] and based on premeasurement essentially of the track vertices only, is now being developed at CERN simultaneously with the operation of the system described above; it uses the same program GATE, but a version of FILTER which takes about twice as long, since it must recognize tracks without any road information to guide it.

Spark chamber film is measured on HPD1 with a different set of programs, similar to the programs for the other type of flying-spot digitizer, Luciole,[6] where the light spot is provided by a cathode ray tube. A whole frame (35×35 mm^2) can be scanned in about 1 sec without moving the film. Distortions are corrected by the computer program, using information obtained by scanning a calibration grid. Luciole is also directly on line to the CDC 6600, and the existing version is used exclusively for spark chamber film, for which Luciole is about three times faster but only half as precise as HPD1. The computer programs[7] for spark chamber film processing, unlike the bubble chamber programs (for HPD), need no guidance by premeasured information. In order to avoid too much re-programming for each spark chamber experiment, the program is organized in self-contained logical blocks and the control of the measurements appropriate to each experiment is centralized in a few routines. The change from Luciole to HPD1 needs in principle only the change of one logical block. The on-line program locates the sparks, fiducial mark positions and identification information for each frame and stores this information on magnetic tape for the off-line stage of the processing, to which we shall return.

• *Other automatic film digitizers*

Although a great number of events (about 2×10^5 events $= 6 \times 10^5$ pictures in three bubble chamber experiments, and about 3×10^6 events in a dozen spark chamber experiments) have been successfully measured on the three

Figure 3 Bubble chamber picture. The crosses are fiducial marks, and the V-shaped track pairs are the interesting events.

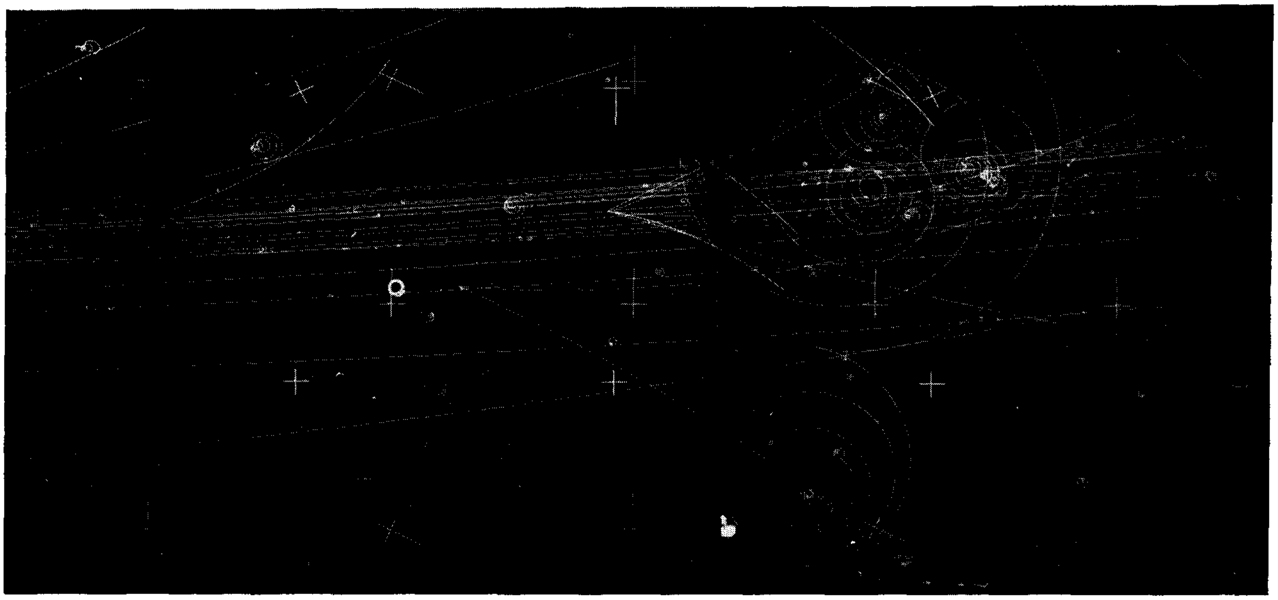

flying-spot digitizers described in the previous section, it has been held against them that the safe, but rather blind and rigorous way they scan the pictures is not the most efficient one. It has also been argued that such devices should not be directly on line to a large batch-processing computer,[8] because of the inconvenient conditions during the development of the devices, and also because the memory they occupy constantly during production is too large with respect to the processor time used. In fact, one spark chamber experiment was measured on a specially made low-precision programmed segment scanner on line to a small computer.[9]

Two new projects exist at CERN. The more advanced is the semi-automatic Spiral Reader, which is practically a copy of the device in use at the Lawrence Radiation Laboratory, Berkeley. It is intended for bubble chamber film, and measuring tests have started. The picture is scanned by a radial slit which moves outwards from the scan center on a spiral path. Each track vertex has to be set by an operator on the scan center with a precision better than the track width, and tracks nearly parallel with the passing slit are digitized. The Spiral Reader is connected on line to a PDP-9.

The other project is an improved version of Luciole (described in the previous section) using a small computer (PDP-9) as an integrated part of its control electronics, which will permit a large autonomy of the system during development and tests. During production, however, it should be on line to a larger computer, preferably one with a fast roll-in roll-out scheme. Under program control it should be possible to address and scan any local region of the frame, and the effective scan area is increased by a mobile film holder. A high precision cathode ray tube should make it possible to measure bubble chamber film. The growing use of filmless spark chambers makes the development of better measuring devices for spark chamber film less necessary.

Computers on-line to electronics experiments

The development of techniques[10] for digitizing directly the coordinates of a particle trajectory has brought, after paper tape, punched cards and magnetic tape, small computers into the electronics experiments. Digital input to the computers comes from filmless spark chambers briefly described earlier under detectors, but also from counter hodoscopes, in which signals from small scintillation counters ("elements") placed systematically in arrays are used for tracing the particles. After the passage of a particle through the hodoscope the signals (or lack of signals) from each element will produce a bit pattern in a register; this bit pattern must later be decoded in order to determine the addresses of the elements which detected a particle.

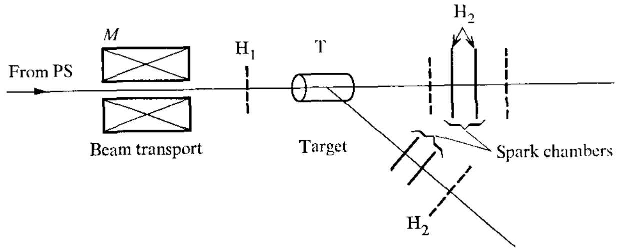

Figure 4 A simple experimental setup.

• *A typical electronics experiment*

A diagram of a simple experiment configuration with analyzing magnets only in the initial beam, is shown in Fig. 4. Protons are injected into the proton synchroton (PS) at low energy, and then are accelerated to high energy and finally ejected; the ejection of the protons (the "burst") lasts typically for 0.2 second, and the time from injection to end of ejection (the cycle) is about 1.5 seconds. The ejected particles pass through a set of magnets and collimators which serves to select particles in a particular momentum range, then through a hodoscope plane (H1), finally striking the target (T) where they interact with the target nuclei. The particles emerging from the reaction pass through further detectors, both hodoscope planes (H2), and spark chambers.

The information from the hodoscope planes—which in simpler experiments may be just single scintillation counters—is used to decide whether or not there has been an interesting event; if so, the spark chambers are triggered to determine more precisely the geometry of the event.

When the chambers have been triggered, the spark positions are digitized and collected in primary storage such as a set of scalers or a small core memory, then transferred to an intermediate storage medium, e.g., magnetic tape, or directly to a computer.

• *Experiments on line*

At present there are about 12 small on-line computers of various configurations incorporated in experiments, involved in the data acquisition, recording and computation of the data. (See, e.g., Ref. 11). None of the small computers is large enough to avoid the need for compromise over the choice of technique of calculation employed. Assuming that the on-line computer is not capable of analyzing completely all events, the choice lies between performing extensive calculations on a sample of events, or limited calculations on all the events. Since in practice limited calculations are little more than instrumental checks, the alternative adopted depends on whether one prefers to have a continuous check of the performance of the experimental equipment, with prompt rejection of obviously bad data, rather than making sure that the data collected have the expected physical meaning.

• *A hodoscope experiment on-line*

The first experiment[12] carried out at CERN on-line to the IBM 360/44—while relatively simple for such a powerful computer—may illustrate the data rate problems that one is likely to encounter in the future; the only detectors used in this experiment were hodoscopes, hence a very high data rate was to be expected.

Preliminary studies showed that the data rate as well as the memory requirement might be critical if the information would be transferred to the computer in the form of the raw binary patterns from the hodoscopes, so a data acquisition system was designed to reduce the data before the transfer to the computer. The configuration of the experiment was such that most of the hodoscopes had less than 64 elements, and in many of them only the case where one particle had passed would be of interest; hence one was able to compress the information from a 64-element (or less) hodoscope—which would normally give a 64-bit pattern—into a byte of 8 bits. Of these 8 bits, 6 were used for the element numbers (0-63), one was used to signal that no element was hit, and one signalled that more elements were hit than thare were bytes provided for reading the data. In the case of a hodoscope of more than 64 elements, or of a hodoscope in which multiple particles were permitted, successive bytes, sufficient in number to hold the required data, were used. In addition to this address information (and some other event information) held in a set of 16 bit scalers, the system held a further set of scalers which were used to record other data transmitted only each PS burst. The system would also operate in a "filtering" mode, where it made decisions itself about whether or not the events should be transferred to the computer, according to the coded information from the hodoscope. In this mode, the second set of scalers described above also held counters detailing the number of events rejected for a particular reason, thus permitting some security. During operation, the first set of scalers were read for every event signalled as good by the system, and the second set once during each PS cycle, at the end of the burst. Encoding, filtering and transfer of one event took about 125 μsec, which was then the maximum effective dead time of the system.

A feature built into this system, which proved extremely useful during setting up of the equipment, and generated a lot of confidence, was the ability to write into it from the computer and read back the patterns, either exactly as written, or encoded in order to perform a check on the encoding circuits.

• *General hardware considerations*

In order to avoid dispersion of effort, one would have liked to adopt a standard system and essentially copy it for all experiments. This has not been done at CERN, partly because successive computers are cheaper and better, partly because about two-thirds of the small computers used for electronics experiments at CERN belong to visiting groups from research centers all over Europe, and finally because data rates and the needs for checking are after all very different from experiment to experiment.

There is still a good number of experiments with data rates low enough to be satisfied with only an IBM 7330 magnetic tape unit on line, but some of these have a very small computer in parallel for instrumental checks.

The most typical small computer configuration has, however, 8-16 K words of memory, a typewriter, two magnetic tape units, a slow card or paper tape reader/punch, and quite often a disk and a display. Interrupts, high speed I/O, cycle stealing channels, memory protests, etc., are, of course, included,

Since these computers serve directly as the fast buffers of the experiments, they are all placed within a few meters of the electronics of the experiments. This is also the most convenient for the experimentalists. It creates, however, some installation problems; because of dust the computers, and particularly the peripheral units, must be in closed huts, which must be well ventilated but small in order that they can be moved easily from one experiment to another (Fig. 5); electrical parasites from adjacent sources are a standing risk.

The only computer substantially larger than the typical size normally installed in the experimental areas is the IBM 360/44 with 32K words of memory and 3 magnetic tape units. It could be used to back up several small computers on line, but it was deemed better to use it directly on line to one large experiment; and if it still has capacity to spare, to make that available off line. If

Figure 5 The installation of the IBM 360/44 in the experimental area. It can be moved easily (e.g., by a crane) in three parts: the air cooling equipment on the roof and the hut with its content in two halves.

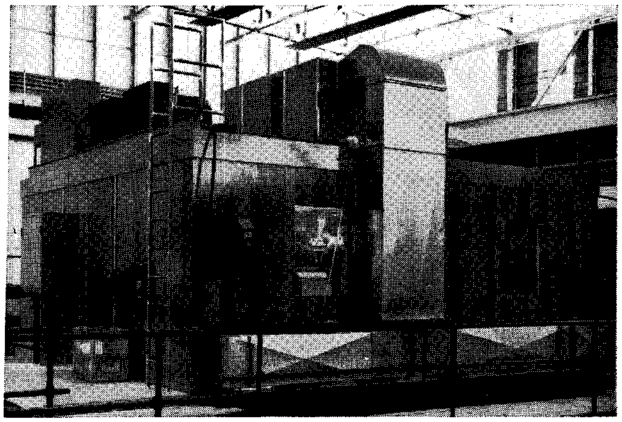

the small computers need to be backed up, they can have a link to the large central computers.

Since in many experiments the raw events can be converted easily into a more economical format, or frequently even rejected as useless after the simplest geometrical checks, another way of "stretching" the capacity of the on-line computers is to perform a maximum of simple logical operations on the data before they enter the computer. This reduces the amount of valuable memory taken up by buffers for useless information and leaves much more memory (and also time) for programs treating the useful data. The system designed for the experiment on line to the 360/44 is an attempt in this direction. Defining the center of a cluster of triggered wires in a wire spark chamber by special hardware is another.

In experiments where the recovery time of wire chambers would be a limiting factor, a further improvement could be to make the triggering criteria more sophisticated. However, with current spark chamber technology these triggering decisions must be made within about 0.5 μsec of the passage of the particle, thus excluding completely the use of present day computers.

Although a graphic display has always been desirable for communication, the cost of memorizing or regenerating equipment has usually been prohibitive. Recently, however, a number of computers have been equipped with inexpensive displays; these are small screen (10 cm by 8 cm) "memoscopes," which memorize the picture on the screen. These have been satisfactory in operation but suffer from two disadvantages—firstly, effects due to lack of precision, which could perhaps be improved by using the larger screens now available, and secondly, the difficulty of providing a dynamic picture, since any change requires that the whole screen be erased and the picture reconstructed.

What one would like is a larger, non-memorizing screen, with a fairly long persistence, and some external decoding circuitry, which would enable the display to be dynamic while using a minimum of core storage, provided that the attached computer had a reasonably sophisticated I/O channel structure (channel multiplexing and autonomous command structure).

• *General software considerations*

Even though many computers have been developed explicitly for on-line use, it is only recently that programing systems have become commercially available to support their special features.

For the first applications of computers on-line at CERN, all the programming was, therefore, done in machine language. For each experiment a system was written consisting usually of data-handling routines from previous systems, and analysis routines tailored for the particular experiment. To some extent this spoiled the experimentalist, for although a huge effort went into each program, the result was usually a program extremely efficient in execution time and space usage. With the arrival of newer computers, with support of on-line applications even extending as far as time-sharing systems, the tendency has been to program in machine language only those operations which would be extremely inefficient in a higher level language, for example handling the input/output and buffering, and to make use of the flexibility of the system and the high level languages for most other tasks. However, in some applications it has been found that the high overheads imposed by these systems, often written for process control applications, have put such restriction on the data-taking capacity as to outweigh the advantages given by the flexibility.

For example, during the last operating period of the accelerator there were three experiments each using small computer of essentially the same configuration on line. One experiment had wire chambers with magnetostrictive read-out, the other two had wire chambers with magnetic core read-out.[3] The magnetostrictive experiment used the manufacturer supplied TSX-system (a time-sharing system permitting sophisticated interrupt handling and core-swapping facilities), since the flexibility and ease of implementation were considered to be of overriding importance; during on-line running it was possible to modify quite easily the FORTRAN-coded analysis programs with no loss of data. One of the other experiments also used this system, but many difficulties were encountered because the core space remaining, after the appropriate system routines were loaded and the data buffers were set up, was inconveniently small. The third experiment ran with a programming system which was built around a much more primitive system. Many routines were reprogrammed and several sophisticated features, such as use of the disk, were added; generality was discarded in favor of minimum space utilization (an over-all gain of about 4K words) and maximum execution speed. The result is a system more efficient than the manufacturer-supplied system for on-line data acquisition, but less convenient to the programmer during the preparation of the experiment, so much more so as the documentation is limited.

In almost all applications of on-line computers at CERN to date, the limiting factor has tended to be lack of memory space; typically one has a program which—with its data buffers—occupies the memory completely, yet may be performing calculations on the data for as little as 20% of the time. While this is an unfortunate mismatch in a small computer used on-line, it becomes economically unacceptable in the larger on-line computers, where it represents a significant loss of computer capacity. A part of the answer seems to be in a simple multiprogramming system, which will permit concurrent operation of one real-time program and one non-real-time program*—the latter could be any operation (e.g. assembly, compilation, execution, etc.)—either by having both programs in memory at the same time, or by having very quick and simple roll-in and roll-out capability.

Currently there is a project under way to "tailor" the programming system of the IBM 360/44 to these requirements for on-line use. A resident data acquisition program will occupy an area in the upper part of the core, providing permanent service for an on-line experiment; at the same time the standard programming system will work as if nothing were changed, except that the effective core size will be reduced by the amount of space occupied by the data acquisition program and buffers, and that one tape unit will not be available. In this way it should be possible to obtain sufficient multiprogramming capability without significantly increasing the overheads of space or time.

Data links to large computers

During the time when an experiment is being set up, and also during production, it is frequently desirable to know the effects of changing conditions on the data acquired. Sometimes the only safe way to see these effects is by thorough analysis of events produced under the conditions, ensuring that a sufficient number of events are analyzed to avoid drawing false conclusions due to statistical fluctuations. The possible saving in accelerator time may justify a very fast access to a powerful computing facility, e.g. by a data link between a small on-line computer or a buffer, and a large remote computer.

The purchase by CERN of a large time-sharing computer system seemed an ideal opportunity to implement a facility of this nature. A link able to transmit a 24-bit word in parallel every 8 μsec was contructed[13] and used during the production runs of an experiment which was on line to a small computer.[14] These production periods showed a data link to be a feasible proposition, but that there were some shortcomings of this particular implementation. This link was, in fact, handled by the CDC 6600 much like the on-line film measuring devices, allocating continuously during the on-line run some 20K words of central memory, which is a very heavy load in a memory-limited computer, even if the average central processor time asked for is low.

A system is now being developed which will overcome many of the deficiencies of this first attempt. This system (FOCUS) proposes to channel and schedule all requests from remote stations for access to the central computers through a small satellite computer. This computer will be the terminal for data links and teletypes, and will permit transmission of data between remote stations and whichever of the central computers is scheduled, accumulation and storage of data, and on-line editing of both program and data files. In this way it is hoped that on-line experiments will be helped during setting-up through having a good access to the large computers with editing facilities for modification of programs, and during production by offering powerful computing facilities on a sampling basis to a number of experiments for up to 24 hours per day.

Data processing

Whether recorded initially on film which is digitized, on paper tape or cards which are converted, or directly by a magnetic tape unit or a computer on-line, practically all the raw data are finally presented event by event on magnetic tape for analysis on the large computers. Since there is little need for random access to the raw data, and the data taking is decentralized and abundant, magnetic tapes are quite suitable, although for reasons of storage space for tapes, overall economy (including operator handling time) and reliability, this could change.

The data analysis programs do not normally contain very sophisticated mathematics, but they may have a complicated logical structure and they should be very efficient, since they are repeated 10^5 to 10^7 times for an experiment.

The first step of the event analysis is to select from among all the information from the detectors or measuring devices events of interest, and to reconstruct their coordinates in three-dimensional space. In film experiments fiducial marks are essential for this. Typical selection criteria might be:

- measured coordinates should form tracks, i.e. lie along reasonable lines;
- the tracks should have an acceptable topology;
- interaction points (vertices) should be in correct regions, e.g., in incident beam or target volume;
- special requirements such as coplanarity of tracks should be satisfied.

In film experiments some of the criteria may have been applied already before or during the measurement, in other experiments in an on-line computer.

Due to measuring errors the selection criteria cannot be applied too strictly, so all events fulfilling them within reasonable tolerances must be accepted.

In order to reduce the effect of errors, redundant measurements and least-square-fits with constraints[15,16] are widely used. Further information on momentum, velocity or range may be provided by (for example) magnets, time-of-flight devices, Cerenkov counters or range measuring arrangements, and a second fit with kinematic constraints[17,18] may permit a more or less complete identification of the particles.

* Editor's note: A Data Acquisition Multiprogramming System (DAMPS) for concurrently real-time and background operation of the IBM System/360 Model 44 is now being made available by IBM.

There is always a danger that the events rejected at some stage of the analysis introduce a bias in the results. A careful study of the nature of these rejected events may be necessary to check criteria and tolerances before the bulk processing, or to recover the events afterwards. For this a CRT display with a light pen has turned out to be very useful.[19] It is easy to localize an error in the criteria if one can see the geometry of the event. In a display it is also possible to eliminate background, and to show the event in a more convenient format than the arrangement of views on spark chamber film (and, of course, in some experiments there is even no film to look at).

The events not rejected during any stage of the analysis are normally stored on a data summary tape, which forms the input to the program giving the final result in the form of a value, formula, curve, histogram, scatter diagram, etc. Among the programs frequently used in this connection are minimization and fitting routines, and the widely used SUMX, capable of sophisticated manipulation and printed graphical output, such as histograms, of any selected part of the data.

Data from bubble chamber and electronics experiments use about equal amounts (about one-third each) of the total available computer time at CERN, but a much larger part of the bubble chamber data than of the electronics experiments data is processed outside CERN, in the member states.

- *Processing of electronics experiments*

For electronics experiments a variety of programs written for each experiment are in use. This is due to the rapid evolution and diversity of experimental technique, to the fact that "tailor made" programs can be smaller and faster, but perhaps also to individualism or lack of communication between groups. However, where experiments have features in common, serious attempts are made to avoid programming duplication. One example of such a feature is the central part of geometrical and kinematical fit programs, which can be taken care of by a package of mathematical routines, whose main task is a series of matrix multiplications and inversions. The key routines were written in assembler language to reduce the overall execution time. To be adapted to a specific case the central package need only be preceded and followed by two experiment-dependent routines.

Another example is iterated tracking of particles through magnetic fields[18,20], which occurs frequently in kinematic fits. A considerable effort has, therefore, gone into producing efficient tracking programs. In many cases it has turned out to be more economic to calculate a certain number of tracks for given momenta beforehand, and derive from those a many-dimensional interpolation formula, giving explicitly the momentum within the required precision. Further work has been done on the related problem of representing the whole field of the magnet economically in the computer[21], and how to obtain this representation without measuring an exaggerated number of magnetic field values. The actual field measurement has partly been automated using a small computer for recording of checked field values on magnetic tape.

Results in the form of histograms are also a very common feature. This is a standard operation with the program SUMX, but many electronics experiments can be well served with a much simpler histogramming routine. To make the use of this simpler routine more flexible a program is being written for the IBM 360/44, which will permit output of intermediate results on printer or display, temporary storage of these results, and modification of parameters by keyboard interaction during execution.

- *Processing of bubble chamber data*

The bubble chamber (or track chamber, TC) experiments have a much more standardized and centralized data analysis procedure, and the description of the available TC program library fills three thick volumes. The most important sections of this library are THRESH (geometry), GRIND[22] (kinematics) and the previously mentioned SUMX. But there are many other sections. For example there is MILLSTONE, a program which permits a close control of the GRIND calculations on nontypical events, although only through special control statement cards in the submitted deck rather than by direct interaction during execution.

Compared with electronics experiments programs the TC library is very stable (but far from static), and the problem of updating and keeping track of all the versions of individual programs in use (necessary in order to achieve generality with economy in computers of different size and make) has led to a reorganization of the TC library.

The programs are now broken down into modular blocks, called patches, and stored in a master file (PAM). A patch may be any suitable part of a complete FORTRAN program. Once included in PAM a patch remains there essentially unchanged, but other patches may be added.

An updating and editing program, PATCHY, steered by appropriate control cards, can reference any patch in PAM, and combines with or without modification copies of them into a complete FORTRAN program, subsequently compiled in the normal way. A patch can also contain control statements and corrections which apply to patches referenced later.

Any version of the TC program can thus be obtained with relatively few control statements, and the same control statements will always produce the same version.

General scientific computing

As the increasing data-taking rate lowers the statistical error in the experiments, the need for accurate knowledge of instrumental quantities is sharpened. For example, the correct and unbiased interpretation may require a careful evaluation of the over-all detection efficiency (e.g., with Monte Carlo methods), or of multiple Coulomb scattering effects.

Accurate experimental results in high-energy physics have, however, only a limited interest if they cannot be compared with an equally accurate theory, or rather, if they cannot be used to check the accurately calculated consequences of a proposed theory. Typical calculations in this category are reaction cross sections, phase space integrals, phase shift analysis, etc. The borderline between this kind of work and the final steps of the data processing is not always sharp.

Before there can be an experiment there must be a good beam. Hence accelerator and beam transport studies, radiation shielding calculations, etc., constitute a further important group of computing applications.

However, we shall leave the subject of general scientific computing here, since it is fundamentally the same everywhere.

Process control

There are other fields connected with high-energy physics in which the presence of computers is making itself felt. One of these fields is process control: accelerators, and certain types of detectors are in many ways ideal subjects for application of some type of computer control, even if this is only data logging and prediction.

At present at CERN, there are three applications of computers of a process-control type. A computer (IBM 1800) is being incorporated into the control system for CERN's 28 GeV proton-synchrotron; and another (PDP-9) is in the control system for the 2 m bubble chamber; it is proposed to use a computer to aid and supervise the operation of the intersecting storage rings (ISR) currently under construction.

Where the computers are being installed after the equipment has been built and in operation for some time, their object is more to relieve the operating staff of some of the more tedious checking, rather than to take over aspects of the operation of the equipment. In these cases the first phase of operation will be concerned with data logging and error detection; following this will be the prediction of appropriate parameter values by interpolation using data previously logged for required operating conditions; subsequently, after satisfactory implementation of these phases, installation of control equipment will permit automatic correction of some error conditions and perhaps modification of operating conditions.

In the case of the ISR application, there will be such a large amount of monitoring and control equipment that manual operation will be next to impossible; for the supervisory system a digital computer, rather than special purpose electronics, is being considered as an integral part of the control system.

As an example, there will be about 300 power circuits concerned with the guiding and focusing of the proton beams, all of which must be precisely and uniformly adjustable. Because of the high degree of dependence on the computer this will be a duplex system; this gives the advantage that basic control functions can be carried out by one or other of the systems, while the second system is available as a standby in the event of an error, or to devote its full power to computation associated with control of the system.

Acknowledgments

The authors are indebted to a large number of individuals at CERN who have given us information for or checked various parts of this paper.

References

1. *Annual Report 1967,* CERN, Geneva, 1968.
2. *CERN Courrier* **7,** No. 9 (1967).
3. W. Baker, *Physics Today* **20,** 41 (1967).
4. M. Benot et al., *Report CERN 68-4,* February 1968.
5. M. Ferran et al., *Proceedings of the 1967 International Conference on Programming for Flying Spot Devices,* Munich, 1967, p. 346.
6. H. Anders et al., *Proceedings of the International Symposium on Nuclear Electronics in Paris,* 25–27 November 1963, European Nuclear Energy Agency, O.E.C.D., Paris, 1964, p. 349.
7. P. Zanella, *Proceedings of the 1966 International Conference on Instrumentation for High Energy Physics at Stanford,* I.U.P.A.P. and U. S. Atomic Energy Commission, p. 507.
8. D. H. Lord, *Proceedings of the 1967 International Conference on Programming for Flying Spot Devices,* Munich, 1967, p. 147.
9. P. Scharff-Hansen et al., *Nucl. Instr. and Meth.* **52,** 25 (1967).
10. *Proceedings of the Informal Meeting on Filmless Spark Chamber Techniques and Associated Computer Use,* Report CERN 64–30, June 1964.
11. G. Bellettini et al., *Report CERN 65-33,* October 1965.
12. R. T. Bell et al., *Proceedings of the International Symposium on Nuclear Electronics,* Paris, 10–13 September 1968 (to be published).
13. F. Marciano, *Report CERN 65-40,* December 1965.
14. T. R. Bell et al., *Report CERN 67-31,* December 1967.
15. R. Böck, *Report CERN 60-7,* March 1960.
16. B. Zacharov, *Nucl. Instr. and Meth.* **32,** 163 (1965).
17. R. Böck, *Report CERN 61-29,* October 1961.
18. B. Zacharov, *Nucl. Instr. and Meth.* **33,** 136 (1965).
19. P. M. Blackall et al, Paper presented at the Computer Graphics Symposium, Brunel University, London, July 29–30, 1968.
20. L. Griffiths et al., *Report CERN 66-17,* April 1966.
21. H. Wind, *J. Comp. Phys.* **2,** 274 (1968).
22. R. Böck, *Instrumentation for High-Energy Physics,* North-Holland Publishing Company, Amsterdam, 1963, p. 435.

Received August 16, 1968