

Data storage

The growth of the CERN library of magnetic tapes over the past six years. There is no sign that the demand for data storage will ease off in coming years and alternative methods of storage are being studied.

E. McIntosh

The workload of the CERN central computer system is dominated, in terms of number of jobs, by short jobs submitted from terminals and remote batch stations. The arrival of the CDC 7600, which will be even more easily accessible and will give the user his results more quickly, is likely to accentuate this domination. Nevertheless, more than 80 % of the system's capacity (whether measured in terms of computation, input/output, or memory utilization) is devoted to the processing of the much smaller number of jobs which constitute the production work of the central system.

This production work consists mainly of the processing and analysis of experimental data by the performance of geometric, kinematic and statistical calculations. It involves extensive computation and considerable input/output from/to magnetic tape and it must be performed regularly, and reliably, to meet the requirements of the CERN physicists.

Recent developments in computer technology, as exemplified in the 7600, should give a significant improvement in the handling of this work. The 7600 is expected to compute about five times faster than its predecessor, the 6600; the new large core memory and fast disks, together with improved operating system techniques, should largely resolve the problem of memory utilization. However, even the fastest magnetic tape units of the 7600 system will be not more than twice as fast as those of the 6600. In addition, the amount of input/output which can be performed by a magnetic tape storage system is limited as much by the speed with which a reel of tape can be retrieved from the tape library and mounted, as by the speed of the data transfer itself.

The increased capacity and reduced access times of the 7600 system


disks will be used to try to ease this problem but the capacity of even the largest disk amounts to no more than the equivalent of a few tens of reels of magnetic tape. This would not satisfy the needs of even individual groups from the many users.

The growing volume of data

The principal source of the data handled at CERN is of course the physics programme at the proton synchrotron, intersecting storage rings and synchro-cyclotron. If we look into the near future in bubble chamber physics — by 1974, when the 3.7 m chamber BEBC and the associated measuring system ERASME are in operation, it should be possible to process over 10^6 events/year. The data collected in this way (raw measurements from LSD, filtered measurements from HPD, or geometry output from ERASME and RAMSES) will then be submitted to the central system for further analysis.

These 10^6 events/year will produce some 5×10^{10} bits of data each year (500 tapes - one 800 bpi 7-track tape holding about 10^8 bits) and this data may need to be retained for up to three years. The analysis of this data at the central system to produce data summary tapes (DSTs) is expected to give at least another 5×10^{10} bits, which must be retained for periods of up to six years for subsequent statistical analysis.

Electronics experiments, in general, generate less data about each event but collect many more events. It is estimated that last year some fifteen experiments yielded 5×10^{11} bits (5000 tapes) of raw data for subsequent analysis. The move towards higher statistics experiments, the availability of beams of higher intensity, and of faster detectors and on-line computers (coupled with the



growth in the number of experiments now that the ISR is coming fully into use) imply that these numbers may be expected to increase considerably this year. It is already possible to envisage experiments measuring events at rates of several hundred per second and producing thousands of magnetic tapes.


As with bubble chamber data, this data will be analysed at the central computing system to produce DSTs, though in this case, thankfully, they should in general have a volume of less than 25 % of the input data. Again, it is usual to keep some or all of the data for lengthy periods of time. Confronted with these figures, perhaps all groups should be compelled to make a vow

'Yea, from the table of my memory, I'll wipe away all trivial fond records.' ('Hamlet' Act I Scene V, as quoted by D. Knuth in his book 'The Art of Computer Programming'.)


1. Photograph (full-scale) of a film chip as used in the IBM 1360 Photo-Digital Storage System. It can contain about 5×10^6 bits (1/20 of a magnetic tape — five copies of Hamlet). The chips are automatically developed after recording and placed in the chipstore. From a maximum capacity chipstore (about 2×10^{12} bits) any chip may be retrieved entirely automatically within a few seconds.

2. Other data storage systems use video tape (an IVC-1000 unit is shown). A reel of video tape can contain between 5×10^{10} and 10×10^{10} bits (500 to 1000 magnetic tapes). The reel may be automatically positioned by a very high speed search to any desired point in an average time of from 10 to 100 seconds depending on the length of the tape.


3. A film strip as used in the UNICON 690-212 laser mass memory system. The strip when fully written by a laser can hold 2.5×10^9 bits (25 magnetic tapes). Any strip from a maximum store of 10^{12} bits may be retrieved in a maximum of 10 seconds.



1.



2.



3.

The limitations of magnetic tape

Magnetic tape handling is probably the worst single source of concern to computer users. Problems may arise with reliability, compatibility of tape units on different computers and transfer of tapes written in the format of one system to a different system. Magnetic tapes are also far from ideal as archives of data; they must be exercised (unwound and rewound) or even copied at regular intervals. However, they remain the cheapest and most convenient medium for recording large volumes of data and for transfers between experiments and computer systems both inside and outside CERN.

Just how popular and cheap it is may be judged from the graph which shows how the tape library at CERN has grown to over 35 000 tapes and is increasing at a rate of over 10 000 a year. At the central system it is

becoming increasingly difficult to retrieve and mount magnetic tapes fast enough to satisfy the system demands and, as stated earlier, there will be a widening gap between the speed at which information can be processed by the 7600 and the speed at which it can be supplied from tape.

The growth, or at least the rate of growth, of the tape library and of tape handling must be limited in some way.

New storage systems

Nuclear physics research is not the only source of large volumes of data. NASA is currently receiving 10^{13} bits (100 000 tapes) of data each year from satellites in orbit; the NASA tape library contains approximately 750 000 tapes. The National Archives and Records Service of the Federal Government of the USA archives approximately 10^{12} bits (10 000 tapes) of

new data each year. In this case fast access is not considered essential!

Many large computer installations (both scientific and commercial) are looking for increased capacity storage systems to replace or complement large disk or tape based storage facilities. The American Government, which is buying about one million tapes each year, is implementing a centralised storage system which will supply archive storage facilities, fast access to a large volume of data, and a convenient means of data communication, to several separate large computer systems.

The first device, based on photo-digital recording on film chip, to be capable of storing on-line 10^{12} bits (10 000 tapes worth) was announced by IBM in 1966. Today there are several devices (including using video recording techniques on magnetic tape or laser recording techniques on film strips) which offer access within

Interactive computing

seconds to on-line stores of 10^{12} bits or more and offer the capability of moving data between the device and an off-line library of reels or strips. Research into the use of holographic techniques has started and could eventually offer much faster access to even larger volumes of data.

CERN is investigating such devices and their connection to its computers; at the same time a data communication network, to make it possible to link the many computers on the CERN site to the central computing and storage systems, is being studied. In this way it is hoped to supply cheap and convenient access to the large volumes of data produced by the experimental programme.

When the first 10^{12} bits store was announced, J.-J. Servan-Schreiber wrote in 'Le Défi Américain'... 'It is estimated that the entire store of information in the world's libraries amounts to 10^{15} (one quadrillion, or one million billion) bits. This information is stored in the form of books and other printed documents, and is doubling every 15 to 20 years..... There appears the possibility that by 1980 a small number of computers will replace *all the written documentation* existing in the world, and that they will work in 'real time' — replying to questions with information at the speed of human conversation.'

Perhaps we can imagine the day when a physicist, from a computer terminal in his own country will have access to any information stored at CERN whether it be experimental results, physics papers or books.

During the past few years a new phrase, 'man-machine interaction', has become fashionable in computer literature. The phrase may be new but the concept it describes is not. Even the simplest computing devices, from the old abacus to modern desk calculators, are essentially interactive computing systems. (Interactive : reciprocally active ; active upon or influencing each other - Oxford English Dictionary). They are interactive in the sense that man is providing the data and specifying the procedure to be followed, and the machine is providing the results while the dialogue between the man and the machine is continuous.

However, with the advent of the electronic computer the pace became faster and the scope more extensive. Particularly when batch processing techniques were introduced, the user became remote from the computer. He precodes his procedures, submits them with the related data and has no possibility of following or interfering with the subsequent process. Only when he bicycles across to pick up his results does he see manoeuvres the computer could have done to get closer to the desired output.

This is not necessarily a bad thing. Overall it can save computer time — it gives long intervals for thought before asking the computer to come into action again. But, sometimes, close and immediate interaction is desirable and several methods of bringing this about are being tried. In particular, there has recently been development of several on-line (also called time-sharing) computing systems, often with graphical output terminals, where the user has a window into the computer and can see it at work directly.

Interactive systems (like relativity theories)) can be divided into two kinds — special and general. Special purpose systems are those where

only a limited class of problems can be treated by the user. More precisely, the user is not able to define a procedure when interacting with the computer, he simply responds to more or less complicated questions from the computer, which has been programmed in advance for solving a special problem. General purpose systems are those where the user is also able to define his own procedure on-line, thus widening the class of problems which can be treated.

This second class of systems can be further subdivided depending on the degree of interaction allowed. Semi-interactive systems are those where the user is able to define a procedure, provide data, manipulate information, and send this to the computer for execution. The user is then not able to interact with the subsequent execution. Fully interactive systems are those in which this latter capability is also present ; the user can, if he wishes, follow the execution step by step, observe intermediate results, change procedures depending on the feed-back, and restart the execution at the point where he interrupted it.

It might seem, at first sight, that a fully interactive general purpose system would be the best solution for any kind of problem. This is not necessarily true because generality will usually be much more expensive than speciality and the degree of interaction should be tied to the needs and, obviously, to the available budget. In the following three sections on interactive computing at CERN we will cover a general purpose system (GAMMA), a file-handling system (FOCUS) with a network of terminals some of which are available for semi-interactive computing, and two special purpose programs where interactive computing (on an ARGUS at the central computers or using FOCUS) has been a great help.