"Record Protection in an Uncertain World"

WORLD CONFERENCE ON RECORDS AND GENEALOGICAL SEMINAR

Salt Lake City, Utah, U.S.A.
5-8 August 1969

STORAGE DEVICES - THEIR POTENTIAL AND RELIABILITY

Part II
Digital Photo Devices

By Garret Boer, B. A.

Part II
Digital Photo Devices

By Garret Boer, B. A.

BACKGROUND AND HISTORY

The Lawrence Radiation Laboratory in Livermore, California, is operated by the University of California under the auspices of the Atomic Energy Commission. It is devoted chiefly to the development of atomic weapons and secondarily to a variety of other research projects most of which require the use of high speed computers. To support this activity, the Laboratory maintains a very large computing facility.

Several years ago, in the hope of improving the service of the computing center, a contract was let to IBM to build a machine which could store a great deal of information and make it available in a very short time. The result of IBM's effort was the IBM-1360 photo-digital storage device. This device can store more than one-hundred million characters and can make portions of this data available in a few seconds. We hope that its use will materially reduce the vast amount of tape storage and tape handling which is now necessary to keep the computers properly supplied with data.

METHOD OF RECORDING.

Notice that the IBM-1360 is a photo-digital storage device. It does not employ a photo-reduction process such as microfilming. Instead, it records basically digital or numerical information by exposing film to an electron beam which moves under the control of data supplied by an associated computer. It is not, of course, difficult to convert alphabetic information to some numeric code, so that records and documents may also be recorded and stored by the 1360. The 1360 must, however, be used in association with a computer which supplies it with the necessary commands to control its operation. In addition, the computer must supply the data which is to be recorded and collect from the 1360 the data which is to be retrieved.

VOLUME OF STORAGE

The basic element on which data is recorded is a small piece of film called a chip - about 35 by 70 millimeters or roughly twice the size of a 35 millimeter slide. One chip can hold nearly five million bits of data or nearly six hundred thousand alphabetic characters if the currently popular encoding of 8 bits per character is used. This is enough encoded alphabetic information to record more than 500 pages of the New Yorker Magazine.

8 bits per char

For storage in the 1360, thirty-two of these film chips are placed in individual slots in a little box called a cell. A cell is about one inch by one-and-a-half inch by three inches, so that you may comfortably grasp in one hand about thirty-two volumes of 500 pages each. The 1360 configuration at Livermore can store nearly 7000 of these cells for rapid access by our computers.

MECHANICAL DESCRIPTION

The 1360 is a large and complex collection of mechanisms which has facilities for automatically exposing, automatically developing, automatically storing and retrieving film chips and finally for automatically recovering and checking data from these chips. A variety of configurations is possible. The one at Livermore occupies about four hundred square feet and has a station for recording and developing, two stations for reading or recovering data, three storage files to hold the cells and a control section to coordinate the device's myriad activities.

The chips and cells are manipulated and transported in a manner which suggests that Rube Goldberg's drawings provided the inspiration. The unit of transportation is the cell, which serves as a bus to carry its thirty-two chips from one place to another; for example, from a storage file to a reading station. This bus is propelled through airtight tubes by differences in air pressure. The transit system uses the same principle of operation that has been employed by department stores to send your money and bill to a central station and to return your change and receipt.

When a cell is to be removed from a storage file, an elaborate mechanical arrangement of trays and tubes is shifted about until one of the 2250 cells in that file is lined up in a pneumatic tube; a burst of air then propels it to its destination, which may be ten or twenty feet away.

When a cell arrives at a reading station - a few seconds later - another sequence of elaborate mechanical actions slide off the lid, remove the desired film chip from its slot and position the chip in the path of an electron beam which will scan it and recover the recorded information.

A chip being recorded is likewise subjected to a complicated sequence of operations. An unexposed film chip is removed from a cell containing raw film and placed in a chamber where it is subjected to a partial vacuum to de-gas the surface. Then as the film is moved into the recording chamber, another chip of raw film follows into the de-gassing section. In the recording chamber the film is exposed to the moving electron beam to record the desired information. After exposure this chip moves out into the developer and the following chip is moved into the recording chamber. Thus, a continuous queue of chips is kept stepping through the recorder.

This queue of chips then moves through a series of stations where they are developed, fixed, washed and dried. They then emerge one at a time and are placed in a cell for storage.

The time required to record and process a chip is about three minutes, but many chips can be in different stages at the same time, so that a finished chip may be produced about every eighteen seconds.

The 1360 is designed to allow many of its various activities to proceed simultaneously. At any one time several cells may be queued up at each of our two reading stations, waiting to deliver their chips, each of the three storage files may be rearranging itself to locate a particular cell position, a cell may be in transit from some point to another and, of course, several cells may be underway in the recording-developing queue.

Control of all this activity is accomplished by the use of a special computer which is an integral part of the 1360. This process control computer must not be confused with the external computer which is still necessary to direct the overall activity. The internal computer interprets commands from the external one. It then sets in motion and keeps track of the dozens of electrically controlled mechanical motions which must be performed in order to obey each command.

1360 ERROR CHECKING AND RELIABILITY

There is, of course, no point in all this activity if the recovered data is not reliable. There are several aspects to the problem of reliability - one is the 1360's internal checks on the data. The 1360 is designed to generate a good deal of redundant information dependant on the pattern of the externally supplied data and to record this redundant information alongside the supplied data. Then, when a chip is retrieved and read, this redundant information is regenerated and compared with that which was previously recorded. This comparison allows the 1360 to determine with a high degree of probability whether the recorded data is correct and even allows the reconstruction of correct data from bad data if the errors are not too severe.

Rather than wait until someone wants to retrieve the data before learning whether or not the 1360 can successfully read a chip, it is good practice to attempt a check-read soon after the chip is recorded - before the external source of the data is discarded. Then it will be possible to immediately re-record a bad chip.

The errors which can be detected internally by the 1360 seem to arise chiefly from small imperfections on the surface of the film which are produced during its development. Our experience at the laboratory so far seems to indicate that we can expect that an average of about one chip per cell - or one in 32 - will have to be re-recorded.

The error-checking described above is strictly an internal check. If incorrect data arrives undetected at the 1360, the redundancy information will be generated from that bad data and when the check is made later, the 1360 will give no indication of error. To further check the data it would be necessary to use the external computer to recover the data from the 1360 soon after recording it and to check it bit by bit against its source.

Once the chip is reliably recorded, it is presumably good forever. We do not yet know whether the chip will indeed deteriorate significantly during a long period of storage.

THE LABORATORY DATA NETWORK

The Lawrence Radiation Laboratory maintains a large computer center. The bulk of the computation is performed currently by a group of powerful Control Data Corporation computers: four CDC 6600's and one CDC 7600. An extensive data network has been set up at the laboratory for at least two purposes.

The first purpose is to connect these computers to inquiry stations. At present there are a couple of hundred of these stations and each one consists of a teletypewriter. We expect the number of stations to grow and we also expect to have more elaborate equipment at each station, for example, a cathode ray tube or TV-like display. From any teletypewriter in the system, a user may communicate with any computer in the system and may initiate some work on that computer.

The second purpose of the data network is to connect each of the basic work-horse computers to the IBM 1360 and to another storage device, the IBM Data Cell, so that data originating at a computer may be stored on one of these devices and the stored data may be retrieved by any of the computers.

At the center of this network is another computer - a Digital Equipment Corporation PDP-6. This computer has data channel connections to each of the larger work-horse computers, to each of the storage devices, and indirectly (via two smaller computers) to each of the teletypes. In addition, the PDP-6 is attached to a high speed disk. This disk is an intermediate storage device which is used to hold data temporarily while it is en route from one place to another.

Before the advent of the IBM 1360 and its incorporation into the network, a computer user had two ways to make large files of data available to his program. First he could enter a deck of cards through the card reader. Secondly, he could type in a message on the teletypewriter. This message would be relayed to a human operator who would in due time seek out and fetch a particular tape reel from our voluminous vaults and hang this tape on a tape unit which is connected to the appropriate computer.

Today, a user may type a message on his teletypewriter which is passed on to the central PDP-6 computer. This computer issues commands to the IBM 1360 and begins the transfer of data from there to the attached high speed disk. When this transfer is complete, the PDP-6 gains the attention of the appropriate peripheral computer, and with its cooperation begins transferring the requested data from the disk to that computer.

The actual data paths are somewhat more devious and a good deal longer than this description indicates. In fact, a teletypewriter message or a data file can follow a route traversing a dozen different electronic devices made by half a dozen different manufacturers.

NETWORK ERROR CHECKING AND RELIABILITY

The lengthy data routes indicated above create additional reliability problems outside of the storage device itself. The reliability checks provided for in the system hardware tend to be very local checks. There is nothing comparable to the elaborate redundancy data incorporated internally in the 1360 which can be used to check the data over the entire transmission path. Some less elaborate redundancy checks are provided at various places along the routes and these do catch many electronic failures and provide a warning to try again. However, we know that many data errors do go undetected by the hardware and that retrieved information is found to be different from that which had been sent to storage.

Programming techniques can be used to a limited extent to overcome this difficulty. Some relatively simple redundancy calculations can at least provide a high probability that you will receive advance notice that some error exists in the data. More elaborate calculations which would allow reconstruction of erroneous sections would require too much computer time to be feasible.

DOCUMENT STORAGE AND RETRIEVAL

The file storage system in use at the Lawrence Radiation Laboratory is not specifically designed to meet the requirements of those who are interested in storing and retrieving documents, and my own work has given me no experience in this application. There are, however, a few general observations I can make about the types of equipment that would be helpful in exploiting a large digital storage device for this purpose.

First, of course, a sufficiently powerful computer must be employed to control the storage device and to handle the flow of information in both directions. There should also be connected to the computer an auxiliary storage device such as a disk or a drum. This auxiliary device will have a good deal less capacity for storing data than the 1360, but could make its data available much more rapidly. It will be useful for at least three purposes: First, it can hold part or all of the index information which is necessary to locate documents on the larger storage device; secondly, it can temporarily hold data which has been retrieved from the 1360, so that documents may be printed from this source while additional retrievals are in progress; and thirdly, it can hold data which has just been recorded until the read-back check is complete, so that the data is available for re-recording if the chip is bad.

One essential part of such a system will, of course, be a card reader, so that documents may be punched on IBM cards for entry into the computer.

There are several considerations which govern the selection of equipment to aid a user of the system in searching for and retrieving a document from some inquiry station. First, it may be desirable to have several inquiry stations so that many people may be simultaneously occupied in searching. A keyboard of some type, such as that on a teletypewriter, is necessary for entering requests into the computer. Also, some searches may require a good deal of conversation with the computer; that is, some initial information is returned by the computer and this prompts the requestor to ask for another retrieval to bring him nearer to the ultimate object of his search. For applications like this, a teletypewriter may be adequate for both entering the request and for printing the computer's response. If, however, large portions of material must be scanned before making another request, a teletypewriter might seem frustratingly slow. For these applications, a device which can display information more rapidly is desirable at each inquiry station. A Cathode Ray tube or television display would adequately meet this need. The system user could scan a few paragraphs and enter a new request via the keyboard.

The TV tube might not, however, be suitable for a more relaxed or lengthy perusal of the retrieved document. Besides, a printed copy is usually wanted for subsequent reference or for use somewhere outside the computer room. To meet this need, a single high speed printer might be tied to the computer to produce this hard copy for all the users of the system.

CONCLUSION:

The advent of the IBM 1360 has moved us into the era when massive collections of data can be speedily accessed. This report is meant to provide a little insight into its workings and some of its problems. I have only hinted at the extent of the programming effort necessary to exploit such a device. It should provide many interesting challenges to those who wish to make use of it.