RICE COMPUTER III

The Rice Computer is located on the second floor of Abercrombie Lab. On the ground floor is a computation laboratory equipped with an IBM 7040 installation, which provides computing service to meet the general needs of the Campus. Several smaller computers for special purposes are situated in other laboratories at Rice. Several large, general purpose computers of various types are accessible within a mile or less of the Campus. Service bureaus are prepared to provide access to computers near and far by dataphone, at moderate rates, directly from any office. No scientist needs to defer his work for the lack of an available computer.

A decade ago, several scientists were obliged to defer their work at Rice because there was not a large scale, high speed computer nearer than a thousand miles or so. One of these scientists was Professor John Kilpatrick of the Department of Chemistry. Professor Kilpatrick had been using the MANIAC at Los Alamos Scientific Laboratory each summer since its completion in 1952. He, along with Professor Salsburg and others, appreciated the need for a powerful computer on the Rice campus. Funds were sought for the purpose, but there were never enough to be found. Professor Kilpatrick continued his efforts and finally, in 1956, secured a commitment totaling approximately one-half million dollars from the Atomic Energy Commission, The Shell Oil Company, and The Humble Company. This was not sufficient to buy a computer, but it was estimated to be enough to produce a replica of the Los Alamos MANIAC II, which had been constructed on site by a small group of engineers there.

Dr. Martin Graham, of the Brookhaven National Laboratory, was invited to take charge of the production. Dr. Graham was a brilliant electronic designer, and came forward with a plan much more ambitious than a mere replica of the existing MANIAC. He assumed his post in September, 1957. After two years, the major logic sections of the system were operable, and Professor Kilpatrick was able to code a geometric puzzle problem called pentominos. Solutions were produced in the form of maps constructed by the computer's high speed line printer.

The computer was placed in full scale operation for regular service to users on Campus early in 1961, and continues to serve to the present day. The computer contained many innovations in its design, and when it was finished, the Rice Computer Project was recognized as having a useful capability in computer research as well as service. A new contract was initiated with the AEC in 1960 to support research activities.

Among the innovations were multiple-level indirect addressing with index register modification at each level, a compound instruction word with several operation fields and several address fields, a set of fast registers for working storage, a pair of data tags for identification and control, and a large selection of indicator registers for interactive control of the computer.

Let us digress a moment to consider some of these innovations.

Data stored in a computer's memory is referred to by a location number or "address." In most machines, the actual address of an element of data must be specified in an instruction in order to operate on that element. However, in sophisticated machines such as the Rice Computer, provision is made to refer to these elements indirectly, hence the term "indirect addressing." In its simplest form, indirect addressing is used to refer

REPRINTED FROM THE RICE ENGINEER

to the data elements through a table of actual addresses (the symbol table). More generally, by use of index registers and special hardware features an element of an array (such as a "row vector") may be accessed by referring through the symbol table. Stored in this table on the Rice Computer, in addition to a base address for the array, is a pattern of bits which determines which index register or register will be used, and whether further indirect addressing shall occur.

The contents of the appropriate index registers are added to the base address, and the resulting address is then referenced by the machine. In the event that more indirect addressing is to occur, this process is repeated until the end of the chain is reached. Hence, for an array of N dimensions, N levels of indirect addressing are needed. Separate index registers are used at each level to specify the appropriate subscript.

Thus the complicated algorithms usually associated with subscripted array addressing have been reduced to a relatively simple hardware feature, affording considerable saving in program running time.

The instructions on all computers consist of various combinations of operation codes and addresses. For example, the familiar IBM 1620 utilizes an instruction format with a single operation code and two addresses. On the other hand, the IBM 7040 is a single address machine with other options in the instruction word. In a single address machine, performing a binary operation (like addition, multiplication, or "logical and") requires the loading of one "operand" (a data element to be processed) into some register with one instruction, and then operating on the contents of that register with a second instruction. Storage of the result of the operation requires the execution of yet another instruction.

On the Rice Computer, it is possible to perform all three of these operations with a single instruction. Thus, what would normally require six references to the memory unit (three data and three instruction accesses) now needs only four (three data and one instruction access). An increase in program efficiency of one-third is thus obtained, if the memory is the limiting factor of the machine speed. This is the case on many commercial computers.

The operation codes on the Rice Computer are divided into eight groups or classes. As originally constructed. two of these performed no function.

However, through the years several specialized instructions have been incorporated into these groups. Probably the most useful of these is analog-to-digital conversion. Through this feature, any one of sixteen inputs can be sampled and stored in the computer as a binary number. The converter will resolve a sample to eight bit accuracy in four microseconds. The actual sampling rate is under program control, and the converter usually operates at ten thousand samples per second or less.

The inverse of this operation is also available. The inherently simpler nature of digital-to-analog conversion makes it possible to provide a continuous conversion of the main and secondary arithmetic registers of the machine.

Another special feature is the ability of the hardware to actuate relays, and to cause time delays in operation which allow these relays to become fully energized or deenergized. This feature can be used to control devices external to th main computer, such as increasing the intensity of an oscilloscope beam to write a spot on its storage surface, or to initiate movement of paper in a strip chart recorder.

By simultaneously utilizing D-to-A conversion and the relay control functions, one may plot alphabetic characters on an oscilloscope face or produce strip chart plots of processed data.

The computer contains many other hardware specialties. In general, the computer is readily adaptable to experimental uses which involve some intervention into the hardware to obtain special operations, special displays, etc.

Innovation at the Rice Computer Project has by no means been limited to circuitry and logic. There are also provisions in the programming systems (software) of the Rice Computer that pioneered more efficient operation of EDP machines. Mr. John Iliffe, now of International Computers and Tabulators, Ltd., England, and Mrs. Jane Jodeit, now of the University of Chicago, were responsible for this work. These innovations have been incorporated in widely-known projects such as STRESS and ICES, (a programming language being developed at MIT to solve civil engineering problems). The Rice system was cited in the software development for PROJECT MAC.

Perhaps the most useful and progressive of these innovations is the dynamic storage allocation. Essentially all commercially available programming systems up until the advent of PL/1 (for the IBM System/360) allocated available memory space to the required data arrays at the time that the source language (e.g., FORTRAN and COBOL) was translated into machine language. This requires that the size of these arrays be known at the time the program is written. In the event that these sizes are not known at compilation time, wastefully large blocks of storage must be reserved to guarantee space for all data. This results in extremely inefficient memory utilization. On large problems, it is quite possible that large amounts of data would have to be stored on relatively low speed memory devices such as magnetic disc or tape. Such problems also inevitably entail considerable manipulation of this data. Computing time may be thus increased far beyond what would be required if an efficient storage system were used. Such an efficient system is dynamic storage allocation.

Dynamic storage provides for the allocation and, more importantly, the reallocation of the available high speed memory space while a proram is running. All that must be specified in the source program is the existance of an array and its name. Note that the DIMENSION statements for FORTRAN are not needed. During the execution of the program, any array may be created, erased, enlarged or reduced as needed. Any storage no longer needed by one array is free to be used to create or enlarge any other array.

This is a particularly important feature for time-sharing and multi-programmed systems. For such systems to operate efficiently, the memory must have the capacity to contain two or more programs and associated data simultaneously. Clearly, if storage is utilized most efficiently, a minimum of time will be spent moving information in and out of the high-speed memory. Since a number of time-sharing systems now operating spend a majority of their system overhead time in deciding what data should be moved in or out of the memory, the efficiency of storage organization is indeed significant. In fact, for many computing problems, the decisions regarding what to store in fast memory may not arise, because all data for the program that is being run at a given time may be held in memory and transferred out only after the problem is completed.

* * * * *

From the time that the Computer was completed, it served the needs of the Faculty and students an average of eighteen hours a day. In busy seasons, it was in use continuously. Operation of the computer and preparation of programs were, in general, handled by the users themselves. Members of the project were primarily concerned with the research program. This involved the development of tapes to accommodate the Geology Department's digital seismograph, the incorporation of cathrode ray tube displays, a recording camera with machine control of film advance, elaboration of the short address fields to perform a new indirect addressing function which could transform abbreviated addresses into full addresses, and many others.

Much of the service aspect of the Rice Computer was taken over by the IBM 7040, after it was placed in service in 1965. However, there are several problems which remain because they are peculiarly suited to being run on the Rice Computer. One in particular, proposed by Professor Salsburg, was deemed to be beyond the capacity of any available commercial system.

Dr. Salsburg's work entails investigation of a problem in statistical mechanics which arose during the study of the properties of solids at high temperatures. It was formulated as a multiple integral of polynomials which were expressed as a double summation. Because of the complicated restrictions on the problem and the complex form of the polynomials, the actual storage requirements for each case were impossible to estimate in advance. Using a non-dynamic storage allocation system would require a memory of extravagant size.

This problem was coded at the Bell Telephone Laboratories, using the ALPAC system, which provided a limited form of dynamic storage allocation. For all but minor cases, the problem exhausted the memory capacity of thirty - two thousand words, and the attempt was abandoned. The same problem was coded in the Rice Computer system, which has a maximum of twenty-four thousand words available. Over a period of time, all cases were successfully completed and the study was completed.

Another problem which was particularly suited to the hardware of the Rice Computer was the "English Spoken Digit Data Sampler," a project carried out by the Rice Systems Group, under the direction of Dr. H. L. Resnikoff of the Rice Mathematitcs Department, and G. A. Sitton of the Rice University Computer Project.

Their work involved the use of the analog-to-digital converter and the data processing facilities of the Computer. A decimal digit, spoken into a microphone, was transformed into binary information in the computer memory by the analog-to-digital converter. Selected properties of the speech, such as mean pitch, intensity, and the spectrogram, were analyzed to determine characteristic patterns.

One of the practical aspects of the project involves speech processing equipment to aid those persons who are deaf because of hearing loss above, for example, one thousand cycles. With the information gained through investigation of typical voice patterns, some sort of translator could be designed to make intelligible sounds which would otherwise be heard only as a "mush" because of the lack of high frequencies to bring out their transient and impulse-like characteristics.

Other computerized operations (such as digital filtering) can be carried out on the stored digitized speech, and the result can be "played back" on the digital-to-analog converter. When speech is played back directly from memory, the fidelity is remarkably good. This indicates that the sampling rate to which the A-to-D converter is limited is sufficiently high to include the vital components of speech.

Research in computer theory and technology is still going on at the Rice Computer. The development of completely new and modernized components (such as the integrated circuit) has given rise to modification of the Computer which will eventually result in the replacement of many of the existing sub-systems.

The most recent addition to the hardware of the machine is a magnetic disc storage unit, which was a gift of the Univac division of Sperry-Rand.

The present memory consists of twenty-four thousand words of magnetic core storage, plus four magnetic tape drives. Each tape drive has a capacity on a single tape of thirty-two core loads of data. But the access time of tape is far from desireable for high speed systems. In fact, it may take on the order of minutes for a drive unit to get from one end of a tape to the other to fetch data. The magnetic disc has an access time of one hundred sixty milliseconds. The discs only have to rotate a maximum of once to make any block of data available to the read-heads. Moreover the disc-file has a capacity of over five hundred core loads of data.

However, adding additional capacity is more than a matter of uncrating a unit and plugging in. Fresh out of the box, the disc-file is imcompatible with the existing memory. The Computer Project staff has designed and is building an interface and control unit to operate between the core and the disc-file. It is interesting to note that since the disc unit operates as a part of memory rather than a part of the computing hardware, it will be compatible with the machine regardless of what other changes are made to the rest of the system.

All new construction related to the computer, including the controls for the disc file, are made from the fastest, most modern integrated circuits available, to postpone the inevitable obsolescence in this rapidly changing technology. In order to deal comfortably with the physical structure problems which arise with the small size of integrated circuits, a standard structural system, known as the Omnicomb, has been adopted.

The physical layout of circuits in this system is determined by a computer program. This program finds an optimum ordering of elements to keep the number of conductors and their average length to a minumum. The program publishes lists of connections, and produces a map of the physical layout, thereby enforcing full documentation of the new circuitry.

Some of the Omnicomb material must be punched to a specification. The optimizing program generates a tape which is used to control an automatic punch which punches the material to the specification.

* * * * *

The Computer staff is always looking toward the future. Software specialists are working on such sophisticated programming sub-systems as "An Incremental Interpretive Formula and Function Evaluator," essentially a system of direct operator interaction with SPIREL (the heart of the Rice Computer operating system) to solve "simple" arithmetic problems without implicit program compilation.

On the drawing boards is a "new generation" of Rice Computer which will have multiple-console and time-sharing capabilities. This system will allow the Rice Computer to be accessed from remote control consoles located around campus, thus making the powerful central processing facilities readily available to the Rice academic community.

Research done with the aid of the Rice Computer has resulted in some sixty published papers in widely separated fields. The staff is always ready to approach new problems with the machine, because it is by this means that the innovations can be put to the decisive test.

* * * * *

The Computer Project has been the proving ground for many students who were attracted by the challenge of computer science and who have gained wide recognition. The March, 1959, issue of *The Rice Engineer*, contained an article by Jane Griffin, entitled "The Rice Institute Computer." Following her marriage to Max Jodeit, her work on dynamic storage allocation in collaboration with John Iliffe became known worldwide.

Mary Shaw, now working toward a Ph. D. in computer science at Carnegie Tech, was associated with the Project as an undergraduate student in mathematics. With Professor Frank Hole of the Rice Anthropology Department, she coauthored an article entitled "Computer Analysis of Chronogical Seriation." This was a pioneering application of the Computer and it was performed at Rice with programs prepared by Miss Shaw. The paper was published as Rice University Studies, Vol. 53, No. 3, Summer 1967.

Joel Cyprus wrote an outstanding thesis on "Optimal Synthesis of the Boolean Functions of Four Variables with Majority Logic." This was published as Rice University Studies, Vol. 50, No. 2, Spring 1964. Mr. Cyprus is a computer designer with Texas Instruments and is a lecturer in electrical engineering at Rice.

As a graduate student, Ernest Sibert implemented most of the works of J. Alan Robinson in Automatic Deduction (the use of computers in theorem proving). This work is known worldwide. Mr. Sibert wrote an excellent thesis on the subject, which is being prepared for publication. Upon receiving his Doctorate, he accepted an appointment in the Rice Electrical Engineering Faculty as Assistant Professor of Computer Science.

A recognized academic program in computer science is gradually being organized through the Department of Electrical Engineering and the newly organized Department of Mathematical Sciences. The Rice Computer Project can function as the experimental arm of this program, and as such it is an asset which not many other universities possess.