A Faster Cratchit

THE HISTORY OF COMPUTING AT MICHIGAN

RESEARCH NEWS

DIVISION OF RESEARCH DEVELOPMENT AND ADMINISTRATION UNIVERSITY OF MICHIGAN

JANUARY 1976 VOL. XXVII NO. 1

Some features of the automatic computer were invented by an Englishman named Charles Babbage in the nineteenth century. Babbage's Difference Engine and Analytical Engine were never brought to perfection. As Charles Dickens has helped to show us, the need for computers in Babbage's time was not evident. Clerks like Bob Cratchit were cheap to hire and adequate to the computational needs of the day. Babbage's engines have been estimated to work at about the rate of a tireless clerk. The speed that made computers attractive was to come one hundred years after Babbage in the form of electronic machines. ENIAC, the first large-scale electronic computer, was not only tireless and able to work continuous shifts, it was thousands of times faster than clerks with desk calculators. Over one Memorial Day weekend ENIAC computed the value of *pi* to 2,000 decimal places, or 1,965 places more than the mathematician Ludolf (1540-1610) was able to achieve in a lifetime. Computers have gotten ever faster over the last quarter century and have been set to performing more and more kinds of tasks. They have greatly altered the pace of scientific research and the style of education in many fields. At the University of Michigan computers have done much for the academic community, and the faculty has made important contributions to computing.

Writer
Blanchard Hiatt
Designer
Sally Marty-Everhardus
Consultants
Robert C. F. Bartels
A. G. Norman
Printed by
University of Michigan
Printing Services

Printed on recycled paper.

The First Computer

When wars have begun and new cannons come into use, gunners have had to be shown how to aim their weapons. In modern times, firing tables have been the means of presenting such instructions to gunners. Firing tables tell where to point a certain weapon to hit a target of a given distance, at a given altitude, in a given wind and temperature. It takes an enormous amount of calculating to produce a firing table because many factors influence the flight of a shell and make its path complex and asymmetric. The shock wave created by its early travel at supersonic speeds must be accounted for, for example, as must the resistance of the air, which differs at different heights and different points along the trajectory. Differential equations are used to sort out the variables for any one set of conditions and to pinpoint where a shell will land.

In 1941 a reservist, Lt. Herman H. Goldstine, then on the mathematics faculty of the University of Michigan, was assigned to the Army Ballistic Research Laboratory at Aberdeen, Maryland, and put in charge of all ballistics computations. He sought a way of hurrying the creation of firing tables. It was largely through the perseverance and inventiveness of a team with which he was associated that the first of the modern computers in this country was created.

The first high-speed, electronic, digital computer was not the invention of any one person. The first in this country resulted from an academic research project sponsored by the War Department. Indirect University of Michigan involvement in the project was strong. In addition to Goldstine, another U-M scholar, Arthur Burks, philosopher and electronics engineer, worked on and gets much credit for the logical design and circuitry of the early ENIAC. Burks is today a professor in the U-M Department of Philosophy; Goldstine became an executive, now a research fellow, at IBM.

ENIAC, electronic numerical integrator and calculator, was the first large-scale electronic computer, but it was not a *stored-program* computer of the kind we know today. Yet ENIAC was extremely successful as a fast electronic means of computing firing tables. By changing the wiring among its various components, that is, by programming it, the machine had the capacity to solve differential equations and arrive at ballistics trajectories and thus took that first major step towards the general purpose computer of today.

Goldstine himself was persuaded of the merits of electronic computing by John Mauchly, who was later responsible in part for the UNIVAC computers. Mauchly had visited a professor and inventor at Iowa State University, John V. Atanasoff, and borrowed some crucial ideas that were to make computers work. Goldstine also saw that the talent needed to build a computer was present at the Moore School of Electrical Engineering, part of the University of Pennsylvania. J. Presper Eckert, later to be Mauchly's partner in the UNIVAC venture, was at the Moore School, as was Arthur Burks. With the war on the horizon, Burks had turned from philosophy to electrical engineering when academic positions were hard to find. He took a government course in electrical engineering at the Moore School and then was invited to join the faculty there.

The production of firing tables was a bottleneck in equipping troops with advanced artillery. At that time firing tables were calculated by hand by small armies of clerks, all women during wartime, using desk calculators and following laboriously a procedure required to account for the many variables. At the end of a hard day, one clerk might have filled in one square in one firing table and begun to work on a second set of firing conditions. There was no available method of appreciably speeding the process. But when Goldstine outlined to War Department officials, including Oswald Veblen, the suggestions for a fast electronic computing machine, a contract was issued to the Moore School to "engage in research and experimental work in connection with the development of an electronic numerical integrator and computer."

Completed in 1946, ENIAC made little contribution to the war effort. It did, however, bring in a new era in computation. ENIAC was at least five hundred times as fast as any non-electronic method of computation, so much faster that it could not fail to overcome any competition. There did exist at the time electro-mechanical computers, but ENIAC with its vacuum tubes operated at a rate unattainable in any machine that employed electrical relays. A relay is a small bar of iron that either opens or closes a circuit in response to a magnetic pull. The electron tube could open or close circuits up to one thousand times faster than the relay because it had no inertia to overcome. It used electronic, not mechanical, means to make a circuit.

ENIAC was one hundred feet long and ten feet high, arrayed in its room like an open-ended rectangle. It held 18,000 electron tubes, 70,000 resistors, 10,000 capacitors, and 6,000 hand switches. It was programmable, and it worked. ENIAC is today in pieces displayed around the country, a part of it now housed in a room next to Burks' office in the Frieze Building of the University.

When the war ended, the ENIAC team went its separate ways, some members, like Eckert and Mauchly, into industry, others into academia. Burks, before coming back to the University of Michigan, spent some time at the Institute for Advanced Studies at Princeton University, where he worked with the renowned John von Neumann. Burks had met von Neumann when the latter joined the ENIAC team as a consultant, and Burks, Goldstine, and von Neumann worked together on the logical design of computers.

At Michigan Burks has continued his work in computers. His studies here follow a course that involves him only incidentally in the nuts and bolts of computers. Rather it aims at elucidating the theory of automatic machines—automata—and comparing their workings to natural (human and animal) automata. His scholarship is the theoretical complement to the sophisticated exploitation of computers that the rest of this *Research News* chronicles.

In 1949 the U-M Logic of Computers Group coalesced around Burks. The theory of automata, which is the focus of the Group's work, is concerned with the ways in which *inputs* to a system combine with *current states* of a system to produce *outputs* from the system. An early goal of the Group was to understand and symbolize a language that would enable one to move from mathematically expressed logical chains to diagrams of the logic and thence from diagrams to actual circuitry. Some of this effort was conducted under a grant from the Burroughs Corporation and was aimed at planning logical designs that use the fewest elements. The Group created such designs, as well as proofs that these designs were indeed the simplest possible. Other aspects of its work include the analysis of self-organizing systems, parenthesis-free notation, and heuristic procedures for investigating cellular spaces.

The importance of this kind of research was attested in 1957 by the creation of an academic program in computer and communication sciences, one of the first such programs in the world, which became a regular academic department in the College of Literature, Science, and the Arts in 1965. Some of the courses offered in the Department are: Man as an Information Processing System, Logic and Automata, Theory of Natural Language Structure, Artificial Intelligence, Advanced Systems Programming, and Software Architecture. In recent years the Department has enrolled about 100 undergraduate majors and sixty graduate students at any one time.

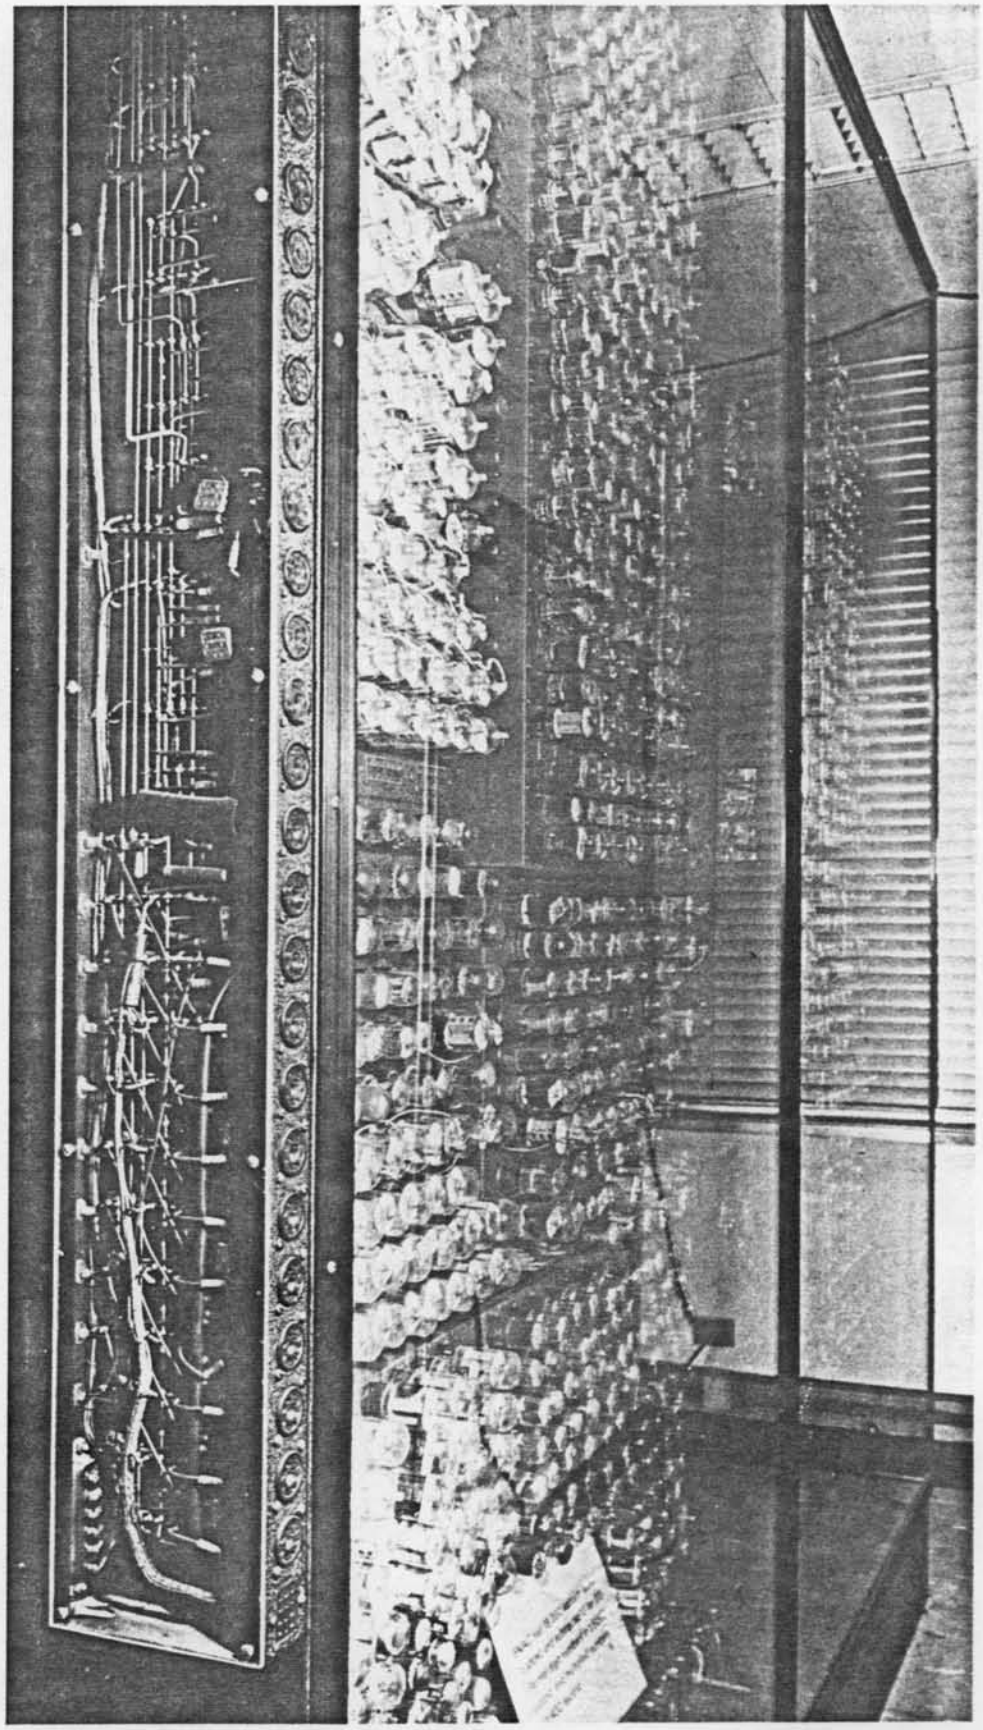

Now enclosed in a glass case and displayed in a room in the U-M Frieze Building, this section of ENIAC (electronic numerical integrator and calculator) suggests how the 18,000-vacuum tube machine looked when it began operating in 1946.

Age of Steam

The era of the steam-driven computer—this is how sophisticates today refer to the 1950's, when automatic electronic computation was just getting started. The University of Michigan was even then active in producing early versions of computers and in acquiring the latest machines that the market made available. In the late 1940's U-M electronics specialists began to build a digital machine modeled after, and improving upon, what was then the most advanced computer. The National Bureau of Standards had designed and built for itself SEAC, which in a local version was to become MIDAC, the Michigan Digital Automatic Computer. It was the -AC period in computing nomenclature: UNIVAC, EASIAC (a programming procedure devised for MIDAC), JONNIAC (a Princeton University creation named to honor John von Neumann), and, inevitably, MANIAC (a computer built by and for the Los Alamos Scientific Laboratory).

MIDAC was in its day one of twenty electronic computers in the United States and one of two in the Midwest. Its purpose was to help U-M scientists working at the Willow Run Research Center solve the thirty-five simultaneous differential equations that described the control of Bomarc, a ramjet-powered anti-aircraft missile of the late fifties. MIDAC was reported in a 1952 issue of the Engineering Research Institute *News* (predecessor of the *Research News*) to be able to "solve problems some 20,000 times faster than a professional mathematician using a desk calculator." Its designers sought to avoid the use of vacuum tubes, the reliability of which was at the time perhaps the greatest weakness in computer operation. In MIDAC, circuitry substituted for many tubes; its 1,000 tubes were only one-tenth the number employed in some of its contemporaries. Also around 1950, a second similar computer was constructed at the Willow Run Research Center for handling certain air defense problems; its development was classified.

By 1954 MIDAC was put to work solving industrial as well as military problems. Operating eighty hours each week, its capacities were by that time enlarged by the use of an early magnetic drum storage unit. The precursor of modern disc storage, the magnetic drum held bits of information in the form of magnetized or unmagnetized spots on its metal surface. The spots could be put on or taken off the drum as it spun around, exposing each spot to a magnetic head with every revolution. Subroutines (built-in programs) and other auxiliary systems also helped to give MIDAC the functions, if not the look, of modern digital computers and to make possible its use by persons who were not specialists in the arcana of machine language.

There were also a number of analog computers on the campus during the 1950's. Analog computers differ fundamentally from digital computers and are useful mainly as simulators of complex physical systems. An analog computer is wired so that voltages flowing through it affect all parts of the circuitry in response to an alteration in one part of the circuitry. For example, the design of an airplane can be simulated (or we may say, an analog of the airplane can be devised electronically) in such a way that one can tell how a change in rudder area, say, affects flight stability, maneuverability, or fuel consumption. The Aeronautical Engineering Department in 1948 was one of the earliest users of a home-built analog computer called the Electronic Differential Analyzer. Other U-M units with early analog computers were the Electrical Engineering Department, the Civil Engineering Department which used its machine to study structural stresses, and the Chemical and Metallurgical Engineering Department.

We do not hear as much today about analog computers as we used to, but this is not because they have gone out of use. They are now more often referred to as simulators than as analog computers, and they are also often linked to digital computers to enhance their utility. Such linking results in hybrid computers that are much faster and more efficient at producing certain kinds of information than digital computers (*Research News,* August/September 1972, "The Simulation Center").

In the last two decades, the greatest growth in automatic computing has been in digital computers, which have overshadowed analog computers in the public mind. Today the word computer, while it still applies to analog devices, has come to be understood generally as digital computation of the sort that calculates paychecks and bills, stores information, and solves problems.

Early campus digital computers, like MIDAC, were few in number, but other types of digital equipment were also present. In 1952 the Statistical Research Laboratory (part of the Graduate School) operated punch-card machines, and the University's Tabulating Service acquired similar gear along with an IBM Card-Programmed Calculator. The Tabulating Service was not part of the academic structure of the University but was rather an administrative service. Now under the direction of Lyle A. Baack, it is today called the Data Systems Center, and it operates a sophisticated computer of its own that is used for administrative and business purposes and not for academic or research computing. Until about the mid-fifties, however, the Tabulating Service did serve research computing from time to time; for example, records from tests of the Salk polio vaccine were analyzed there.

This practice of using record-keeping tools for both administrative purposes and academic research was to change soon, and wisely so. One of the key decisions in the development of computing at Michigan was taken during the tenure of Professor Emeritus Ralph A. Sawyer, who was then dean of the Graduate School and, later, vice-president for research as well. Sawyer and others recognized that administrative and research computing were to a degree incompatible. Administrative work would too often relegate research computing to second priority. Each month when the payroll had to be gotten out, all other uses of the computer would have to be suspended. By the sixties file security had become yet another reason why administrative computing had to be kept separate from research and instructional computing. The research computer is accessible to several thousand persons at any time, and although the system is carefully contrived to protect the data that belongs to each user, clever pranksters have demonstrated on occasion that files can be read surreptitiously. Personnel files, held in the Data Systems Center, of course cannot be left vulnerable to intrusion of any sort. Hence the wisdom of keeping entirely separate the two kinds of University computing.

A Proper Computer

Digital computing for research and instruction was thus under way in a variety of units of the University by the mid-1950's. One such unit, the Statistical Research Laboratory, then under the direction of math professor Cecil C. Craig, became relatively prominent when it obtained from IBM in 1955 a Model 650 digital electronic computer. A unique pricing policy established by IBM made it attractive for universities to procure this and subsequent IBM computers. The rental agreement between IBM and the University stated that the University was free to use the 650 for all purposes but that it need only pay for that portion of computer time that was occupied by sponsored research projects. IBM said that it expected about forty percent of the 650's time would be so occupied and that hence the monthly rent would be forty percent of the normal rental rate for the 650. For several years IBM was thus to subsidize the use of its computers both at this University and others.

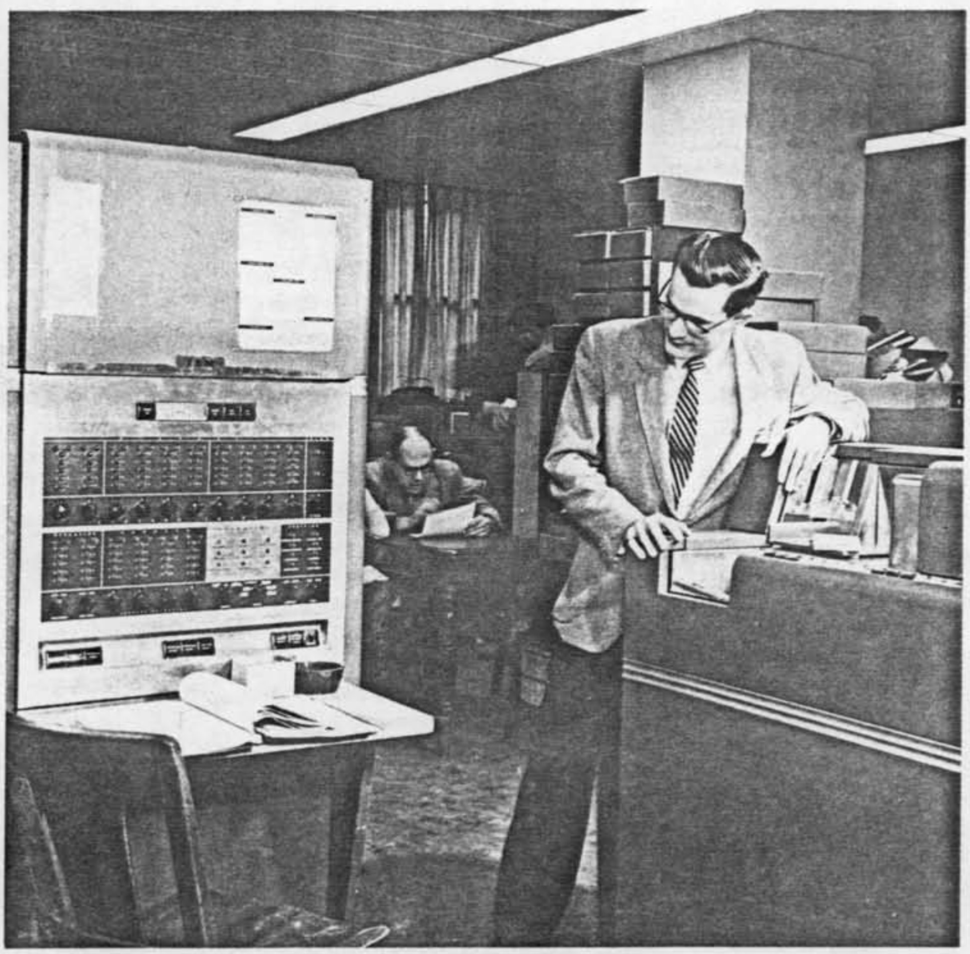

The IBM 650 was the first University of Michigan computer that was generally available for research and instruction. It was housed in the Statistical Research Laboratory. Shown here is the 650's console and a card-reading and card-punching machine.

The 650 was a respectable-sized, reasonably flexible computer, the sort of machine that newspapers of the fifties used to call an "electronic brain." It was more reliable by far than earlier computers and permitted the execution of a series of programs without continual operator intervention. It also allowed for the design of *programming languages* of some sophistication, that is, methods of communicating orders to the machine without having to resort to machine language. Machine languages were different for each machine, in accord with differences in their designs, and were of no interest to anyone but computer specialists. Indeed, the necessity to use machine language hindered most early users of computers because it compelled them to break down their work into elementary logical statements understandable to the machine. Problem-oriented languages, which the 650 made practical, allowed users to cast their computer programs in terms familiar to themselves. A few hardware buttons pushed on the outside of the machine in a certain "grammatical" sequence activated many software buttons on the inside of the machine. And when the program was finished running, the results would be re-translated into the user-oriented language before being printed out.

(Let us take a moment to explain some basic jargon. Computer *hardware* is pieces of equipment: the central processor where logic and arithmetic are handled, the input devices like terminals, the output devices like high-speed line printers, the circuitry that comprises a computer's memory, the tape machines and tapes that hold programs and data, the telephone lines that link user to computer, and other items that can be seen and touched. Only a few kinds of *software* can be touched, for example, manuals and instruction books. Other software is mainly ways of making the machine work, intellectual order imposed on the hardware in the form of *programs*. Many kinds of programs are used time after time or continuously. These kinds of software include the program that manages all other programs on a computer and programs that translate the schematized but human-intelligible problem-oriented languages into machine language. Software keeps a computer in touch with itself and with the outside world of users.)

The 650 was the beginning of routine, large-scale research and instructional computing at the University. Faculty members and students devised a language to make use of the machine's capacities, called GAT (Generalized Algorithmic Translator), which greatly increased the U-M computing population. And the machine's reliability inspired the creation of operating programs that would efficiently manage a series of users' programs, which came to the computer in batch after batch of punched cards.

Officially the 650 worked from eight till five and was unplugged every evening. But its eight-hour day was a full one, with users on hand to exploit every minute. After the 650 had been retired and replacements for it had arrived on campus, stories began to emerge about the earlier machine's double life. After hours, it appears, students gained access to the 650 through a window in the Stat Lab, which occupied (as it still does) the ground floor of the Rackham Building. By this method the number of computer-competent U-M scholars got still another boost.

In part because the 650 was the first versatile computer on the U-M campus, it helped to show that new institutional arrangements were needed to make it accessible to the growing numbers of persons who could benefit from it. Here was a tool capable of solving many kinds of problems and getting easier to use with each passing year. Yet it was part of only one department, indeed a part of a unit within that department, a machine evidently intended mainly to serve mathematicians and statisticians. It was in fact available to all for the asking, but there was some evidence that researchers tended to see it as a tool of the Mathematics Department. New arrangements, to take effect when the next generation of computers came on the market, would help make it plain to all that the computer was available to all.

Ford Foundation

By the late fifties it had become clear that computers were going to become increasingly bigger, faster, more useful, and more common. In fact, computers were developing so fast that it seemed all of society would soon need a crash course on computer programming. One place to begin such a course was in the prominent institutions of higher learning. The Ford Foundation asked the University of Michigan if it would conduct a Project on the Use of Computers in Engineering Education. The University was to operate the Project for the benefit of college teachers of engineering from a wide area around Michigan. In 1959 the Ford Foundation provided $900,000 to pay for equipment, visiting faculty, consultants, postdoctoral students, and conferences.

Up to 1958 some graduate engineering students had had the need or the incentive to elect a non-engineering course on computing, or they had taught themselves programming. This method, while producing some excellent results, lacked the thoroughness needed for a sound foundation for using computers in engineering. To meet this new need, the Ford grant helped the University to assemble a team under the leadership of Donald L. Katz, then chairman of the Department of Chemical and Metallurgical Engineering. Working with Katz were faculty members Norman R. Scott, Robert C. F. Bartels, Glen V. Berg, Stuart W. Churchill, Bernard A. Galler, William P. Graebel, Geza L. Gyorey, Robert M. Howe, Donald R. Mason, George L. West, Franklin H. Westervelt, Dean H. Wilson, and Richard C. Wilson. Professor E. I. Organick of the Computing Laboratory of the University of Houston was appointed assistant director of the Project.

The Ford Project sought to teach teachers how and when to use computers so that they might in turn involve their own students in computing. By 1963 the Project had trained 220 engineering teachers, including more than seventy in the U-M College of Engineering. Participants were eager to learn about computing, and many of them, following their twenty- to thirty-hour intensive courses, immediately introduced computing into their own courses here and at many other universities.

As important as teaching teachers was the production of instructional materials. The Project paid for the preparation of a large number of completely programmed and solved engineering problems and for their publication. These paradigms represented a variety of problems and met the needs of students in all engineering fields. U-M Project members also developed an introductory course in computing techniques and equipment for engineering undergraduates.

By 1958 the use of the computer in research and instruction had begun to affect many programs at the University of Michigan. New institutional arrangements were therefore needed to make the University's digital computer visible and accessible to everyone. Thus in January, 1959, the Regents of the University established a Computing Center. They did this upon the recommendation of a committee that had been chaired by Vice-President William Stirton. The Center was to be a research and service unit of the Graduate School. It would provide equipment and consultation to both researchers and to teachers who wanted to involve their students in computing and would be housed in a remodeled section of the old North University Building.

Robert C. F. Bartels, Professor of Mathematics, was named the first, and to date only, director of the Computing Center. Bartels was respected around the University not merely for his background in computing but for his teaching abilities. "If he believes that I'm the one who persuaded him to take the directorship," one of his colleagues said recently, "then I feel a little guilty, because the Center soon occupied all his time and took him away from the classroom." Bartels' experience with computers dated back to 1954 at Oak Ridge National Laboratory, where he had used an early computer in the study of magnetohydrodynamics (which deals with conductive fluids and their interactions with magnetic fields). His mathematical expertise lay partly in numerical analysis, a field which became increasingly important during the forties and fifties because of its relevance to computing. At Michigan he had worked with student dissertations that involved computers, and he was an active member of the Ford Foundation Project.

The creation of a computing center had several advantages. As a unit under the dean of the Graduate School, it was obviously a facility meant to serve all academic interests. As a central facility, it took upon itself the goal of meeting the various special computing needs of different campus users, so that users were not faced with obtaining their own special gear for their research projects. Through joint appointments the Center established personnel links with academic units; its senior staff members might include at any given time a professor of engineering, a physical scientist, a social scientist—persons from units in which computing was important.

For it is unworthy of excellent men to lose hours like slaves in the labor of calculation which could safely be relegated to anyone else if machines were used.

Gottfried Wilhelm Leibniz

The administration of the Computing Center by Ralph A. Sawyer, dean of the Graduate School, also led to other benefits. It was necessary, Sawyer felt, that computing, if it was to develop rationally on campus, must pay its own way. That is, the Center must get along without subsidy from the University, paying all its bills from funds generated by charges for computing to projects supported by research grants or contracts. Sawyer knew that in those days computers were still looked upon with suspicion. They had not proved to everyone that they were either necessary or economical. Ill-conceived or too rapid "computerization" in business and elsewhere was known sometimes to cause fiscal debacles. If the Regents and the executive officers of the University were eventually to support computing, Sawyer felt, it had better be presented to them as a cost-free operation while it established itself.

Fortunately, IBM was still giving a price break in the form of the so-called "educational allowance," by which the University could in effect use the computer without charge for non-research purposes. Of further help was a three-year grant of $150,000 from the National Science Foundation to establish and operate the Computing Center.

A secondary effect of the self-support policy led to another long-run benefit. Because the Center was not able to afford to hire a staff of programmers, every U-M computer user had to learn to do his own programming. At other institutions this was sometimes not the case. One turned his problem over to a computing staff programmer and hoped for the best. At Michigan each researcher, each teacher, and even each student who used the computer remained responsible for meeting his own programming needs. From the U-M Computing Center one could get only help (limited to counseling on difficult programs or programs gone awry), never hindrance in the form of white-coated technicians who stood between the machine and the user.

Dean Sawyer later became the first U-M vice-president for research, while continuing his role also as dean of the graduate school. When Professor A. Geoffrey Norman succeeded Sawyer as the vice-president for research upon the latter's retirement in 1963, the Computing Center was transferred administratively from the Graduate School to the purview of Norman's office. The computer was now fast becoming an indispensable resource for instruction and research, and Norman was soon to face the problem of acquiring University General Fund support to supplement support that came from sponsored research projects.

UMES and MAD

The progress made in administering computation at the University during the late fifties was complemented by progress in computer technology. Scheduled to be the new Computing Center's first computer was an IBM Model 701. But delivery of this computer was delayed so long that the IBM 704, which superseded it, was finally installed at the Center in August, 1959. This was not the first of the large computers to follow the IBM 650 at the University, for there was also the larger IBM Model 709 at the Willow Run Laboratories. At the Laboratories little, if any, instructional use was made of the IBM 709 computer. Its principal use was for purposes of simulation studies carried on under the University's Project Michigan with the support of the U.S. Army Signal Corps. The continuing financing of the IBM 709 computer became difficult for the Laboratory, for when the funding for Project Michigan was curtailed in 1960, the primary source of funding for the computer was drastically reduced. This resulted in a transfer of the administration of the IBM 709 to the Computing Center in January, 1961, and then to the actual physical replacement of the IBM 704 in August of that year. The IBM 704 and 709 were both rental machines. With the transfer of administration to the Computing Center, the IBM 709 was made generally available for use by the University-supported instructional and research projects and thereby became eligible with the IBM 704 for the rental discount that IBM offered.

IBM's educational allowance program, which was in effect prior to March 30, 1963, made computer equipment available to the colleges and universities at forty percent of the company's published monthly rate. This undoubtedly was a factor that influenced the selection of equipment by many of the schools of higher learning throughout the country, and certainly was among the several determining factors in the selection process for the first line of computers at the University of Michigan. And it undoubtedly hastened the spread of computers and training in computers throughout the colleges and universities. But the benefits were mutual, for the exposure of generations of students to IBM equipment during their formative years as professionals was certainly not unanticipated.

The IBM 704 was ready in the fall of 1959 with the new Computing Center to serve the computational needs of the University community, and the Ford Foundation Project was one of its most active users. The 704 was a much bigger computer than the 650 that had preceded it, and in certain operations it was at least one hundred times faster. So too was the IBM 709, which replaced the 704 in August, 1961. The 709 was, in fact, fundamentally like the 704, but it had connected to it special subsidiary computers, called channels, that managed the components of the system peripheral to the central processor, such as storage and input-output devices. By September, 1962, the 709, too, was to give way to the last of the IBM 70- series of computers, the 7090, which was a transistorized, or solid-state, version of the 709. Designed without vacuum tubes, the 7090 was six times faster than its predecessor.

The potential of the 704, 709, and 7090 computers necessitated the development of software and systems programming that would effectively exploit them. The Computing Center staff came up with two important developments. It devised a new problem-oriented language suited to U-M purposes and to the 704's design. Called MAD (Michigan Algorithmic Decoder), it excelled in efficiency the comparable FORTRAN language provided for the 704 by its manufacturer and was recognized as a major intellectual achievement. MAD was developed for the specific purpose of training large numbers of university students and handling the large volume of university research problems. The main objectives in the design of the language and its compiler were the speed of translation, generality, ease of use, and ease of adding to the language. MAD proved to be a significantly faster translator than FORTRAN. Equally important, it provided features that were advantageous to the programmer and not available in other languages at the time. But certainly because of its speed of translation it was a very useful facility for running a large number

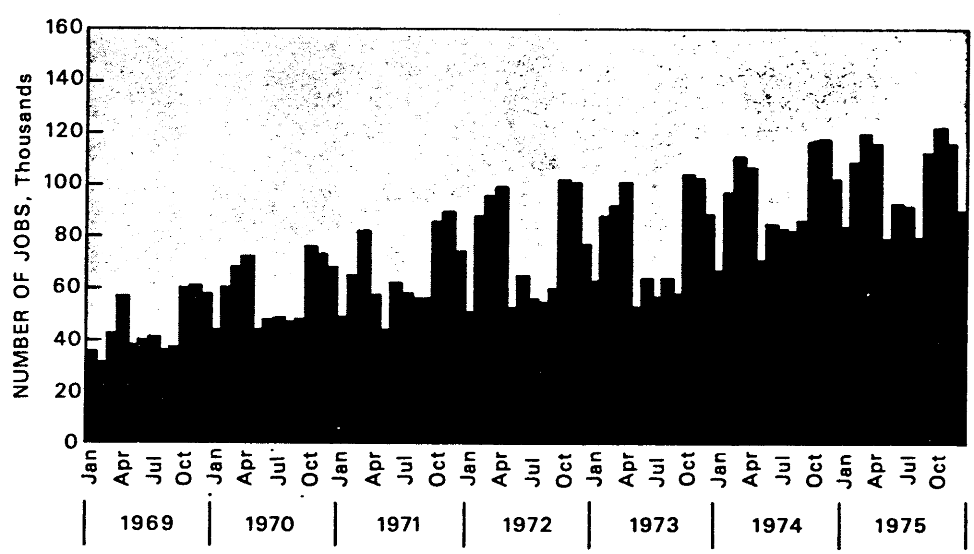

This graph shows the growth of computer use at the University of Michigan.

of student problems on the IBM 704 and later in adapted form on the 709 and 7090. It was adopted by a large number of other users of these machines, including as many as twenty-four other universities. The creators of this language, which has helped make simpler the useful application of the large-scale computer to students and researchers in almost all parts of the world, were Bruce Arden, Bernard Galler, and Robert Graham.

The Computing Center further departed from the use of the manufacturer's software on the 70- series computers by adapting to its own needs an operating system developed at the General Motors Technical Center in Warren, Michigan. The Michigan version was named UMES, the U-M Executive System. (An operating system is a set of rules or conditions under which a computer operates; all programming on a computer has to be compatible with that machine's operating system, just as all plays in a football game must accord with the rules of the game.) The basis for the selection of this executive system and for the changes wrought in it by the Computing Center staff under Galler's leadership was the need to handle efficiently the large number of jobs that characterize a university computing facility. The goal was to minimize the system time for the very common translate-load-execute student jobs without seriously increasing the system time for the fewer but often more complicated production jobs submitted by researchers. Within the constraints of the sequential job-stream automations that these systems emulated, the goal was well met.

UMES was also adopted, like MAD, by other universities and organizations. IBM promoted the spread of UMES by helping to pay for the preparation and distribution of documentation in UMES and for a series of lectures on its structure and operating procedures.

Crisis of Rising Expectations

MAD and UMES were so successful in meeting computing needs that the ironic inevitable happened. By 1963 U-M students and faculty members had begun to saturate the computing system and also had come to expect more than was technically feasible on the equipment available in the Computing Center. Computer applications that were now contemplated by innovative scientists led to demand by U-M researchers and teachers for the direct control of and interaction with their programs during execution on the computer.

To interact with a computer means to sit at a console, usually a teletypewriter or a variation thereof, and type instructions to which the computer responds immediately. Rather than deliver one's program to the computer in a batch of punched cards and then wait to pick up the results, one can in the interactive mode converse, as it were, with the computer via typewriter and telephone line. Interactive, or *on-line,* processing complements the batch mode of computer use.

The Computing Center could not in 1963 provide on-line service to U-M users. This would have meant shutting many users out while one party got all the computer's attention. Huge chunks of processor time would go to waste while that one party sat thinking at his console. The only practical method of operation then at the U-M Computing Center (as almost everywhere else) was to continue taking in batch jobs on punched cards and returning print-out results a few hours later or the next morning.

But some U-M computer users needed more than batch-mode processing. They had heard of early efforts to develop interactive computing, and they knew that they needed it if their work was to progress. By 1963 the Mental Health Research Institute (MHRI) and other units were eager to get their hands on an interactive computer system. In fact, MHRI in 1963 submitted a proposal requesting federal support for a relatively powerful computer aimed at serving an interactive operating system. Such a scheme was to meet the needs both of MHRI and of certain other university researchers. Although the proposal was not approved in its entirety, MHRI did receive grant support for a modest computer to serve real-time, man-machine, interactive requirements of some of its research.

It was at this point that the University's central administration became concerned over the prospect of proliferation of computers on campus. To insure orderly development of a computer resource that would meet all of the diverse needs for instruction and research, Roger Heyns, then vice-president for academic affairs, created the *Ad Hoc* Computer Advisory Committee. Faculty members in a variety of disciplines were selected for this committee, and Dr. Heyns appointed Donald L. Katz, Professor of Chemical Engineering, as its chairman.

Choices were now to be made afresh on how best to protect and promote the development of computing on campus while meeting the needs of the greatest number of users. The Katz Committee was charged not merely with resolving the matter of interactive computing, though this came to be the central issue; it also faced such other problems as the request by some users, notably researchers working on masses of NASA data taken from space flights, to purchase whole shifts of the computer's time for their own use. Other users besides MHRI were also requesting that they be permitted to purchase their own small or medium-sized computers. Research grants were in those days relatively generous, and some users who wanted to avoid waiting in line for the 7090 had money to purchase their own machines. Since there was often considerable merit in such purchases, the Committee had to decide whether each peripheral computer should be re-

A Computer's Day

A computer is so expensive that its human masters take pains to fill every one of its seconds with a billion nanoseconds of computing. The fast-adding servant compels the masters to spend whole careers just thinking up ways to keep it busier. Perfect busy-ness is achievable (not including time-out for maintenance) by batch processing under control of an operating system like UMES (p. 10). But human inconvenience rears its head, and, in spite of the costs associated with idle processor seconds, users want a system like time-sharing. With time-sharing, it becomes inevitable that when users are not lined up waiting for the computer—when they are conveniencing themselves by sleeping or eating or opening a window or attending to the other inefficiencies entailed in being human—during such times a computer must be suffered to pass some idle fractions of seconds.

As a time-sharing system, MTS is designed to meet the needs of users whenever the users call on it. Thus it experiences peaks of use that accord with those times when most users have work on their minds. The Michigan computer's day includes three peaks, one in the morning, one in the afternoon, and one in the evening. Fortunately for MTS, many of its clients are students, an irregular bunch who are content to work late. Daytime work peaks are thus partially mitigated by the choice of a good fraction of MTS users to work hours outside of eight-to-five. A normal day's computing at a normal point in a semester can be graphed as follows.

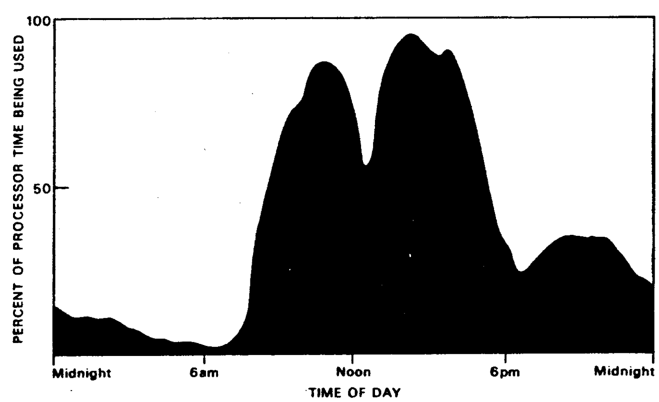

Daytime use rises sharply in the morning, falls somewhat at lunchtime, then goes highest at about 2:30 P.M.. A suppertime lull is followed by a rise in the evening.

At certain times, the computer's day takes on a profile like the following, with flat-topped peaks.

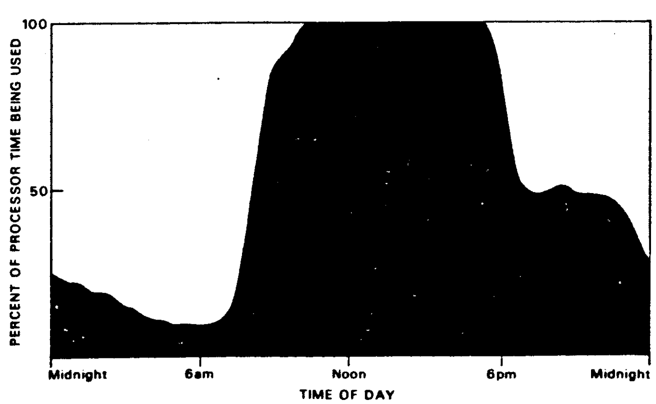

When days that look like this become common it means that MTS is overloaded. The flat peaks show that the processor is at a given hour being used to capacity. Peaks like this can be expected late in each semester when students are paying the price for having convenienced themselves too much in prior weeks. When such patterns prevail, however, it suggests that the University needs a bigger or faster processor. This was the case in 1973 with the 360/67 system.

Let us begin a day with the computer at 8 A.M. and observe its use. The first twenty hours are available for use by anyone with an MTS account number (to pay for his computer use) and a password. Both these must be entered correctly at a terminal before a user can get on-line to the computer. During morning and afternoon peaks of computer use, the processor approaches full use. Only about one second out of ten goes unused. By late evening computing drops off.

At 4 A.M. MTS ceases to answer the phone. The Computing Center now orders the system to perform filesaves for the next two hours. All the programs that have been called up and worked on or put into the system during the previous day have been recorded on magnetic discs. The discs are expensive and have to be erased for re-use each day. Files are saved as the discs' contents are recorded on magnetic tapes, which are cheap to store. Now every user's program is preserved in its current form. Another kind of filesave is performed weekly in the small hours of Sunday: all of the week's programming is then copied on tapes. This duplication of filesaves ensures that programs, the loss of which could sometimes ruin or delay a career, can always be recovered. Periods given over to filesaves also allow time for the Computing Center staff to experiment with MTS and to test new ideas for better service. For this the staff needs command of the entire machine.

At 6 A.M. the computer and all its peripheral gear are turned over to vendors for maintenance. The Computing Center staff does some maintenance, but an important part is conducted under service contracts with the companies that supply the equipment. For a monthly fee the Computing Center is entitled to replacement parts as needed and to the services of a repairman in the supplier's employ. During the early months of the Amdahl 470V/6's service at the Computing Center the supplier had two repairmen on call to the University.

The weekly exception to this schedule occurs between 11:30 P.M. Sunday and 6 A.M. Monday when the entire system is shut down. Holiday periods also cause schedule changes. Like Bob Cratchit, MTS gets Christmas off.

quired to be compatible with the central computer (so that the two could be linked for a long-term advantage) and whether there ought to be a size limit on peripheral machines.

The Katz Committee did its work during 1964. It formed itself into subcommittees representing the different kinds of computing needs felt by workers in different fields. The subcommittees held open meetings to ensure that everyone on campus with an interest in computing might be heard. The open meetings did away with any feeling that decisions might be imposed by fiat or by a cadre of technicians, and the academic community was confident that all decisions would respond to the real needs of users.

Central to the recommendations made in December, 1964, by the Katz Committee was that the University pursue as rapidly as possible the creation of a time-sharing system. In this recommendation the Committee was still ahead of its time. Though time-sharing did exist by 1964, operating systems needed to manage it had not been perfected. Indeed, computing hardware was not yet up to the demands of full-scale time-sharing, as the future was to demonstrate. Nevertheless, the emphasis on time-sharing was appropriate. A restive faculty threatened to move elsewhere, the report intimated, unless the University upgraded its computer services to include interactive, remote-terminal computing managed by a time-sharing system. Other universities, notably the Massachusetts Institute of Technology, were venturing onto the new ground, and so should Michigan.

The Katz Committee also ratified some other concepts and suggested new policies, many of which have been in effect for over a decade. The Committee's report endorsed the existing separation of administrative and scientific computation. The report also endorsed the concept of a central computing facility like the Computing Center, then five years old. A central computer was the best insurance against costly duplication of hardware and the best guarantee of maintaining a base of users who could support the purchase of the biggest and fastest equipment.

At the same time, the report recognized the need for small computers in some units, especially where real-time operations were called for. Real-time operation means that a computer is interacting with only one party during a given interval and usually that the problem under study involves measures of the passage of time itself. Studies of man-machine interactions, for example, and the monitoring of medical patients require constant measures of time and cannot be done unless a computer, often a small one, is dedicated to the one task and not interrupted by any others. There are also other important uses for small computers. But, while endorsing small computers, the report held out against their unlimited adoption and also held that each small computer (except those dedicated to one purpose) should be linked, like any remote terminal, to the central computer. Linking greatly enhances the powers of each small machine. The report called for a panel to review requests for peripheral computers, to counsel on the selection of each, and to prevent any unnecessary proliferation of small computers. The Committee sought to protect the central computer by maintaining a large number of users for it.

Finally the report proposed that the *Ad Hoc* Committee be continued as a Computer Policy Committee. It would advocate on behalf of the *Ad Hoc* Committee's report, advise the vice-president for research, and mediate between faculty and the vice-president's office as necessary. In particular, it would screen requests for small computers and, as it has often done, refine faculty members' plans for computer use.

The Library Model

Administrative changes in computing, as well as systems changes, were a hallmark of the early sixties. In late 1963 Dean Sawyer retired and was succeeded in his office as vice-president for research by A. Geoffrey Norman, Professor of Botany. One of the first matters that Norman had to attend to was the effect of the decision by IBM to modify its rental agreements with educational institutions, making colleges liable under U.S. Bureau of the Budget regulations for their full shares of the cost of their computers. All computer users now had to pay the same rate for computer time if the University was to remain eligible for federal research grants. This meant it was no longer possible to recover the full rental cost from only those projects supported by external sources of funds; rather, it required a sharing of costs by the University in proportion to its own computer use for instruction and internally supported research. This change reflected the mounting pressure on IBM by federal agencies, which were seeking a more equitable sharing of the market for computers and computer services. Unfortunately the new price agreement brought about the end of what in effect had been the free use of the computer by the students and faculty at the University of Michigan.

The change in IBM's agreements with universities was made effective in March, 1963. The U-M policy which had been in force, namely, that the computer be self-supporting and without cost to the General Funds, was at a sudden and involuntary end. It became necessary at once for the University to pay for all computer time used by its students in equivalent dollars from internal sources of money. Since in reality about a third of the computer's time was expended for instructional and non-sponsored faculty research purposes, the burden of finding internal funds to cover this much of the cost of the IBM equipment was suddenly thrust upon the University.

Negotiations with IBM in an attempt to gain respite from sudden inroads into the University finances lasted about a year. Since the University's budget for 1963-1964 had already been released to review by the state legislature long before IBM's announcement, every effort was made to seek an acceptable extension of the former agreement under which the Computing Center had operated, at least for the 1963-1964 fiscal year. The University was not alone among the large universities in suffering from the sudden change in the procedures for computer financing, and it is possible that IBM dared not set a precedent through a special concession to the University of Michigan. Thus, in the spring of 1964, after a year of unsuccessful negotiation, the University accepted the new agreement with IBM and established a procedure for financing the University's computers under the agreement. The task then for the newly appointed vice-president was not only to acquire funds to cover the costs for the "free" use of the computer over the past year, but also to persuade his executive colleagues and the Regents that the contributions of the computer to the education and research programs would justify the drain it would impose on the budgets in the years to come.

The accounting procedure instituted through the efforts of Dr. Norman and the Computing Center persists virtually unchanged to the present. The Computing Center's operations are financed entirely out of the funds derived through charges for services rendered. By a schedule of pricing, the costs over a fiscal period are recovered from the funding sources available to all users in proportion to their use of each of the services derived from the Center. Hence the University now makes available a source of funds to support student and faculty use of the computer for course or University-supported research projects. The procedure for charging research projects which had external sources of funds had already been established and remained unchanged. Of course these projects experienced a substantial reduction in charges owing to the expanded distribution of the billable costs over all users of the Center's services.

The billing procedure was extended to include all users by allocating money from the General Funds (supplemented during a transitional period by a grant from the National Science Foundation to help pay for the use of the computer in the sciences) to support the use for the non-sponsored instructional and research projects. The use of these funds is apportioned among the colleges, schools, and other units by an allocation of quotas in the form of vouchers in amounts that are determined by past use and by projected needs submitted with the annual budget requests by the units.

In short, the University had to move from a posture in which the Computing Center was essentially self-supporting to one in which the Center was expressly budgeted to give service to all potential computer users in the academic community. This accorded with Vice-President Norman's attitude that computation services should have the same institutional role as library services. The cost for the service should be borne by the University budget as a whole, and not by actual charges to individual student users or users without external financing. Like the library, the Computing Center would now provide services to all fields without regard to how well funded they may be by research sponsors. Of course, it took General Fund money to accomplish this, but Norman found that he could persuade his executive colleagues to provide this money in a series of steps. Large-scale computation was, at last, firmly embedded in the academic structure of the University.

Time-Sharing

Thus 1964 saw the *Ad Hoc* Committee's urging of a new level of computer systems sophistication as well as new institutional arrangements for supporting such advanced developments. In the same year the Computing Center staff began in earnest to contemplate its next hardware purchase. Such a new machine would exploit the latest technology and operate fast enough to meet the projected future load of computing; it would also have to be a machine adaptable to time-sharing, that is, a machine suited to complex multiprogramming.

The time-sharing system that everyone hoped for was to be able to handle over one hundred programs at one time. The users were to be able to sit at their terminals, identify themselves, and type instructions to the computer, such as "Fetch my data from my file and find its square roots for me." The computer would then seem to carry out the instructions immediately, printing out or displaying the results at the requesting user's terminal. Owing to the great speed of the system and to the slowness of human activity at the terminal, the results of an instruction would appear so fast that users would feel that they were getting the full attention of the computer. To accomplish this, the operating system would automatically switch the control of a processor over to the terminal users' programs in a round-robin fashion, giving each a turn for a thin slice of time. At the same time, one or more batch jobs, depending on the terminal load, would also be given turns at slices of processor time.

One of the first Computing Center staff members to work on time-sharing was Robert Graham, a research associate with the Computing Center. Graham spent summers in 1962 and 1963 at MIT, bringing home the results of early studies of time-sharing techniques. Graham helped persuade the Computing Center staff that time-sharing would not only prove possible but also inevitable; he then departed for a position at MIT.

Bruce Arden and Bernard Galler were also active in time-sharing plans, and much of the Computing Center's activity in 1964 centered on finding a new computer that would improve on the 7090's capacities and lend itself to a time-sharing operation. One potential replacement for the 7090 was the Control Data model 6600, which then had the world's fastest processor. But a study conducted in the fall of 1964 by an advanced seminar in computing showed that the CD 6600 could not handle more than about thirty terminals and would rapidly prove inadequate to U-M needs. Arden, Bartels, Bernard Galler, and Franklin Westervelt prepared specifications for a time-sharing computer and sought responses from major manufacturers. Several companies offered to supply equipment that would meet most or all of the specifications. A variant of the recently announced IBM System/360 line of computers, which conformed in principle to a scheme of machine organization proposed intially by Bruce Arden, was selected as the best suited for time-sharing.

The decision to once more procure an IBM computer begins one of the most interesting chapters in computing at Michigan. The decision involved the Computing Center staff deeply with IBM in the development of time-sharing and greatly influenced that supplier's hardware and future marketing plans. IBM's principal computer system to this day embodies the U-M influence.

Norman gives special credit to the triumvirate of Michigan computer specialists who at this juncture were to contribute greatly to the future of computing at Michigan and in the nation as a whole. "Bartels, Arden, and Westervelt," Norman has said, "were a team that we took great care should not be broken up or induced to leave the University. Westervelt, the hardware expert, Arden, brilliant in software and logic, and Bartels orchestrating their progress—these three put together a superb time-sharing computer system. The University and their faculty colleagues owe them much."

IBM had at first apparently not planned to include the machine selected by the Computing Center in its line of marketable computers. This computer, labeled the System/360 Model 66M (the M in 66M standing for Michigan), was to be one of a kind built to University of Michigan specifications. Fortunately, the tide of interest in equipment that would allow for efficient time-sharing was on the rise. IBM quickly found that the computer proposed to the University of Michigan in accordance with U-M specifications was attractive to other buyers too. Within months the hardware innovations that were unique to the 360/66M were to become standard features in the newly announced System/360 Model 67 computer. This had significant advantages for the University both in price and maintenance, for within a year of its decision IBM had gathered orders for *forty* 360/67's, and the system took its place as IBM's standard product.

Thus the 360/67 hardware and time-sharing system owe much to the computing aspirations of researchers and teachers at the University of Michigan. The specifications that were drawn up in response to the *Ad Hoc* Committee's recommendations became the basis for a new product that has met the computing needs of many users in the United States and elsewhere for over a decade. Further collaboration on machine hardware and software between IBM and other early customers for the 360/67, in which Bruce Arden assumed principal leadership, deepened the influence of the Computing Center on IBM and, what is important locally, brought to Michigan the most advanced computing system.

The Step into the 360/67 Era

In early 1965 the executive officers approved the expansion of the Computing Center's budget to allow for acquiring the IBM 360/67 and signed a letter of intent to IBM for a computer that had not then been built. Many months later, during the waning days of December, 1966, the computer, then with only a single processor, was delivered to the North University Building. Vice-President Norman had seen to raising funds to modify and expand the building to accommodate the new equipment, which was to share space with the IBM 7090 and, for a short period, with an IBM System/360 Model 50 that had been introduced to supplement the 7090. Arranging equipment around the pillars in the North University Building was never easy, and the presence of three machines congested greatly the available space. Even so, the 360/67 was installed and, on March 7, 1967, accepted by the University.

The 360/67 was initially retained on rent with an option to purchase the machine that could be exercised later. As yet there was no experience with time-sharing or with the new machine. Only time would tell whether the new system would work effectively and, if it did, whether it would be capacious enough to handle the growth in computing that the interactive mode was expected to inspire. The rental plan provided the University with an escape route if the system were to prove disappointing; on the other hand, in case of success, arrangements made it easy to replace the single processor with a duplex processor, thereby doubling capacity.

By hindsight these plans were overly cautious. The system did prove a success and was expanded in August, 1968, by the addition of a second processor. In the spring of 1969, in a move that would ultimately save the University a good deal of money, Norman took advantage of the purchase option and, with consent of the computing community and the Regents, bought outright the duplex 360/67 from IBM. This purchase began an era of software stability and user convenience that is continuing today.

One reason that the University could feel confident of its ability to finance the purchase of the large computer was the National Science Foundation's generous support of U-M facilities expansion. Two proposals, made in December, 1965, and in May, 1969, argued that support in behalf of computation for research and instruction at an institution such as Michigan was in the national interest. The grants totaled $1,180,000 and eased the difficult transitions of the 1967-71 period, when the 360 was installed, when the 7090 was supplanted (in July, 1968), and when computer use grew greatly in many fields.

The 360/67's key characteristic was flexibility: it could serve a variety of types of use, and it could grow through the addition of many kinds of peripheral modules. It also had built into it special circuitry, some of which was a product of the Computing Center's original hardware specifications. This circuitry facilitated the time-sharing operation. The speed and capacity of the 360/67 made possible the expansion of Computing Center services. Measured in internal machine cycle times, the 360/67 was eight times faster than the 7090. As for memory, the 360/67 even with only a single processor had four times the capacity of the 7090.

What the 360/67 computer did not have at the time of its delivery was a functional time-sharing operating system designed for the special features of the machine. IBM had in progress the development and construction of such a system for use by its customers on this computer, but its hardware people outpaced its software staff, and the software system was continually delayed. Here was some firm evidence that the *Ad Hoc* Committee's call for a time-sharing system had been in advance of its time. While IBM struggled with the software to meet the schedule of machine delivery, the Computing Center staff, anticipating the delay, began to fashion its own software. The Computing Center had no choice. There were research projects dependent upon the capabilities of an interactive computer that were waiting for service at the time of delivery of the 360/67. That there was no other choice has since been hailed as good fortune. It led to development of the Michigan Terminal System (MTS), which delivered to the University community the time-sharing interactive system that the Katz Committee had requested and which exploited many of the unique features of the 360/67.

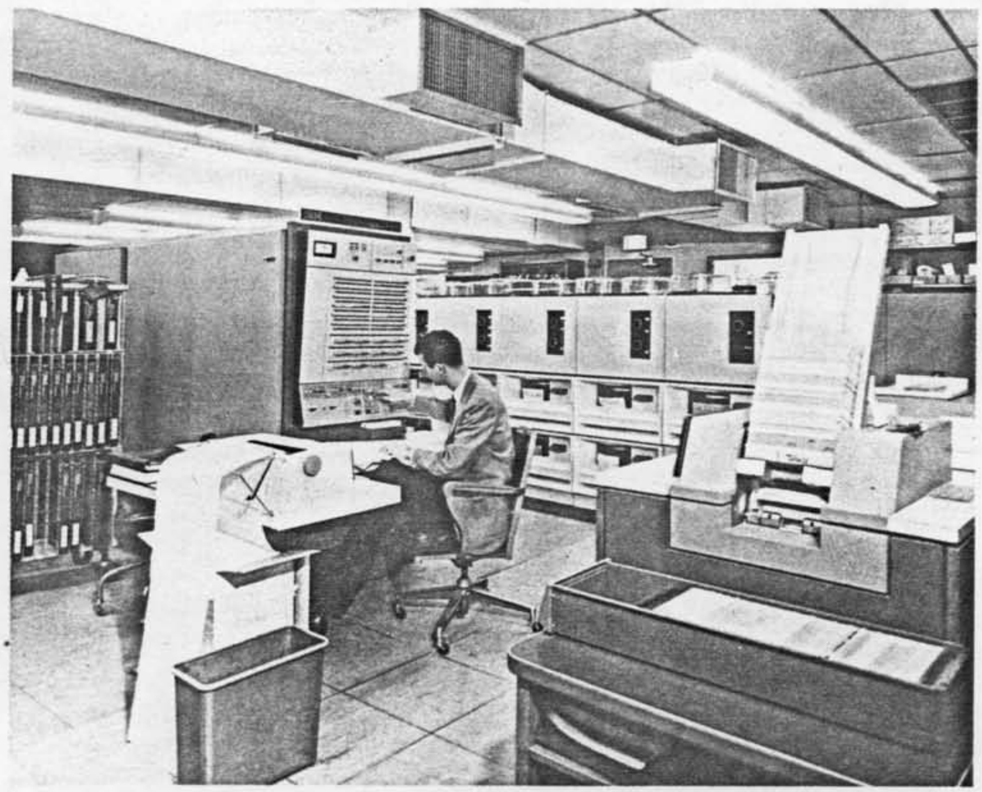

For about half its lifetime on the U-M campus, the IBM 360/67 was housed in the North University Building, which got many modifications to enable it to accommodate computers of increasing scale and to serve a rapidly growing community of computer users. The typewriter-like device pictured here is a user terminal; here any instructions and information may be typed in, and results can be printed out on great ribbons of paper. With such a terminal users can interact directly with the computer, seemingly getting instantaneous response from it, as if one had the processor's entire attention. One could also use the 360/67 in batch mode by means of a card reader (right) under the control of the operating system, Michigan Terminal System (MTS).

Just as MAD and UMES had been major intellectual contributions to computing in the previous generation of hardware, so was MTS to prove a brilliant tool for 360/67 users. There have been many contributors to MTS, but the credit for molding its original structure is shared by Michael T. Alexander and Donald W. Boettner of the U-M Computing Center. MTS was actually to prove itself better able to exploit the machine architecture of the 360/67 than the operating system that IBM would eventually create. This is hardly surprising, since so many of the ideas that influenced that machine's design had come from the same Computing Center staff members who now were devising an operating system to put their ideas to work.

In one way or other the requirements of many of the research projects gave ideas and purpose to the design of certain of the functional characteristics and capabilities of MTS. One of special significance because of its contributions during the period of early development was the University's ConComp Project which was supported by the Advanced Research Projects Agency of the Department of Defense, directed by Franklin Westervelt and Bertram Herzog.

MTS took control of the 360/67 in

November, 1966, eleven months after the machine's installation. It instituted interactive, on-line computing and instantly raised by ten times the Computing Center's capacity to serve U-M users. The key idea that made this possible was the *virtual address,* which is described in a famous paper, "Program and Addressing Structure in a Time-Sharing Environment," by Arden, Galler, Westervelt, and T. C. O'Brien. Virtual addresses are a method by which an operating system like MTS, rather than the computer user, assigns storage locations to the user's program. A program is broken up into many small pieces by MTS and stored at many addresses at MTS's convenience. Yet from the user's point of view, the program as a whole, or any part of it, is available to him virtually instantly. The user no longer has to concern himself with the actual address of his program, which is not even assigned until the moment MTS begins to execute his program.

The virtual address greatly expands the computer's ability to work on several programs in rapid alternation. It thereby makes possible the interactive mode of computing. Terminals scattered across the campus (in 1976 there are about 200 of them) allow users to dial the computer at any time and use their own or some other permitted files. They may execute a program, ask that some data be displayed by the terminal, manipulate their data to get new information, or alter their program. The form of display is commonly typescript printed by typewriter-like devices, but it may also be a television-like cathode ray tube display, or even a synthesized voice speaking through the telephone that attaches most terminals to the computer.

A unique device, created by the Computing Center for the ConComp Project in 1966, enables MTS to link a variety of kinds of terminals to the computer. The *data concentrator* is a kind of butler that answers the computer's door and renders all messages into a common computer-intelligible form before carrying them to the master for processing. One reason for this is that, by analogy, some callers speak with a heavy accent or very slowly, and the master grows impatient waiting for them to finish. That is, telephone line-speeds and other terminal and program features differ, and the computer's processing efficiency may be greatly impaired unless some agent buffers it from the outside world.

The first data concentrator was built around a small peripheral computer, the PDP-8, made by the Data Equipment Corporation. It intercepted all communications with terminals and adjusted signals in both directions to meet the technical needs of both the processor and the user's terminal. The Advanced Research Projects Agency of the Department of Defense helped to pay for the data concentrator's development because the Agency was also interested in methods of coping with line-speed differences.

Conventional batch jobs can also be handled on MTS. A punched card reader is just another kind of input device, like the interactive terminals. It enters new batch programs into the system, and a line printer presents the results. The extensive modification and adaptation in 1969 of an IBM subsystem called HASP, for Houston Automatic Spooling Priority, helped improve batch-mode service on MTS. HASP supervises the queuing of the batch jobs to be processed and the subsequent output services, such as printing and punching, that are required. In an attempt to optimize job through-put, HASP assigns priorities for the ordering of the jobs or tasks within each of the several queues and the rate at which they are released for processing.

MTS operated successfully on the 360/67 for eight years, from 1966 to 1974. In effect it brought a powerful computer into every office, every laboratory, and every teaching unit where computation was needed. With a twist of his finger any user could dial up and use any part of the enormous variety of Computing Center software and other conveniences. The flexibility of the system was widely recognized too. Today MTS is in use as the principal operating system at the University of Newcastle-upon-Tyne, the University of British Columbia, the University of Alberta, and Wayne State University.

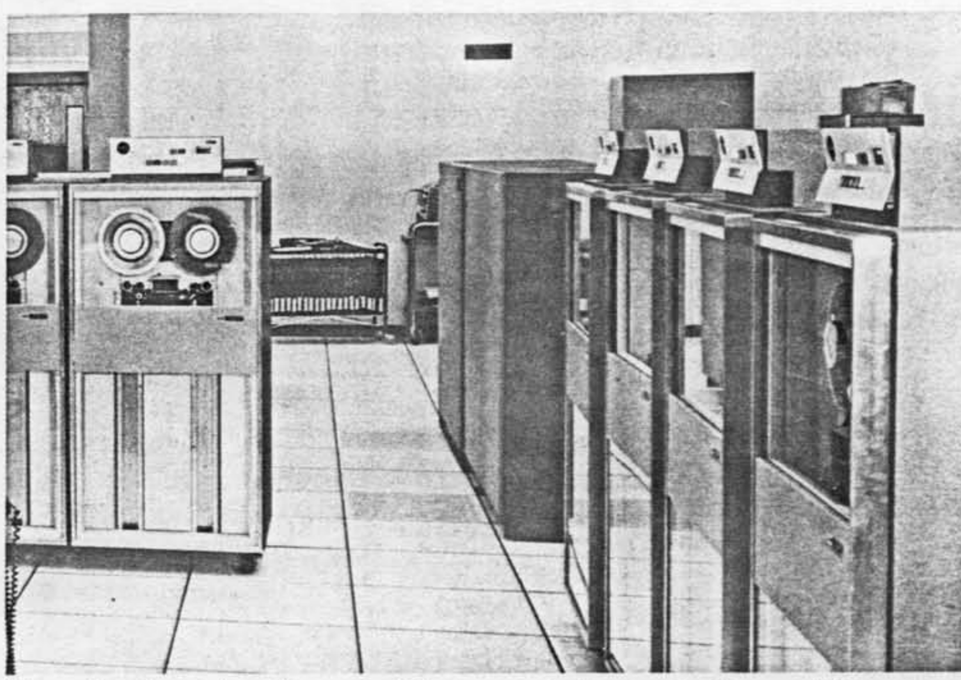

When the Computing Center building was completed on the North Campus in 1971, the IBM 360/67 was moved there from the North University Building and put into operation after a mere three weeks of service interruption. Here one can see the long, reinforced concrete beams that support two floors and the roof. The upright column visible in the background bears no load but is merely a conduit for electrical connections between floors. Computer and gear are stacked on three floors to minimize electrical transmission distances, which can hamper the efficiency of fast, modern computers.

No Pillars, Hollow Floors

"It astonishes me," an MTS user said recently, "that the speed of light has become a consideration in my daily life." For him the universe's speed limit has become a practical matter because it limits the rate of information processing in modern computers. Today it is necessary to design computers and the buildings that house them with a care for the distances that electrical signals must travel between parts of a computer system during processing. The Computing Center is now located in a building that eliminates some architectural constraints that had begun to hamper the efficiency of MTS.

For more years than anyone wished the Computing Center occupied the central campus's North University Building, located between the Museums Building and the Heating Plant. The North University Building at first provided adequate space for the IBM 704 and for services to the limited user population of 1959. But the building quickly became overcrowded. Students, who came to the building to fetch their results, might find their program had failed; they filled every inch of corridor and floor space as they remained in the building to correct their programs and resubmit them.

Three expansions of the North University Building, two of which were aided by funds from the Ford Foundation and the National Science Foundation, created an atmosphere of perennial flux. The University sought during the mid-sixties to raise money for a new building. IBM, approached for a gift, declined to give money directly for a computer building on any one campus, but it did give an un-earmarked gift of $300,000 to the 1967 Sesquicentennial campaign. When in 1969 attempts to persuade the Office of Education to fund a new building proved vain, the University decided to proceed on its own.

Both the timing of the new building's completion and its location were the subject of careful deliberation. Norman aimed at getting the building erected during a period when the 360/67 was continuing to provide service, a period when no equipment needs competed with funds for the building itself. The location of the building on the North Campus was protested by some, but objections faltered in face of plans to maintain a sophisticated input-output facility in the North University Building on the central campus. In this way student users and others without their own terminals were not inconvenienced by the location of computer facilities at a distance from the bulk of users. Terminals are now located in many campus buildings and at the U-M Flint and U-M Dearborn campuses, where students and faculty get their computing done by MTS. Portable terminals are also found around the University especially for use in the classroom; these can be put on line to MTS from any room that has a telephone.

The Computing Center building was ultimately erected at a cost of $1.4 million. In designing a building to house a large computer, there was little experience to draw upon. Moreover, later expanison and dimly perceived technological advances had to be borne in mind. Some of the building's architectural graces deserve mention. Load-bearing outer walls support long spans of reinforced pre-cast concrete beams that carry the upper two stories. Pillars thus cannot constrain the placement of computer and peripheral gear, which are stacked up on each of three floors to minimize transmission distances. User services share space on the lower floor with input-output devices. The second floor holds the computer console, communication control facilities, and disc and tape storage. On the third floor sits the processor, channels, and fast storage devices.

Because computers require a narrow range of temperature and humidity if they are to work dependably, the building has a powerful air-conditioning system, which is rendered virtually fail-safe by backup equipment. The building's cooling capacity would suffice for over forty average houses. The hollow floors serve as a plenum for this large amount of ventilation, eliminating masses of ductwork, and they serve as well as passageways for the miles of cables, which are the sinew of MTS. Nine hundred telephone lines serve the building and make possible future growth in the number of terminals on MTS.

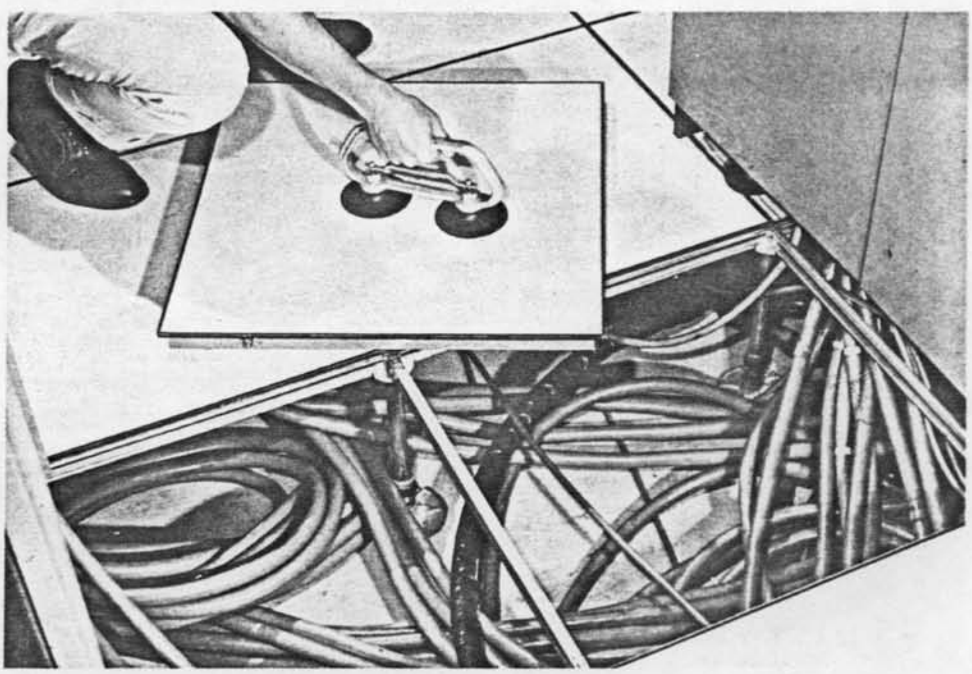

The Computing Center building has hollow floors to which the staff can gain ready access by means of rubber suction cups. The hollow floors not only carry wires but serve as a plenum for the air-conditioning system.

Off-Line Subculture

No computing system, however deftly organized, can meet all computing needs. Real-time operations, noted above, have not been served by MTS. A real-time use to which small computers are especially suited is the manipulating of experiments in split seconds in response to changing parameters; even slight delays, as might be entailed in waiting a moment to get the attention of MTS's processor, are sometimes intolerable in these applications. Small computers can also provide instantaneous analysis of output from certain experiments. Many kinds of investigations are made possible and the rate of scientific advance speeded when experiments can be designed around computer monitors located right in the laboratory.

Thus the Computer Policy Committee has over the years approved several dozens of requests for the purchase of computers for specific applications (not always restricted to real-time). Some of these computers have been large. The computer at the Institute for Social Research (ISR), for example, is the IBM 360/40, managed under IBM's Operating System. A small staff of computer specialists works with the different groups at ISR, and this staff has over the years developed OSIRIS, a group of general purpose programs aimed at easing the analysis of surveys and at other statistical operations. OSIRIS has been adopted by many users who are dependent upon the equipment and operating system of IBM. ISR's Computer Services Facility processes more than 50,000 jobs annually.

The Department of Physics owns a computer, a PDP-10, purchased with a grant from the Atomic Energy Commission. Research in physics, particularly particle physics, calls for massive statistical operations. The large volume of this specialized programming, nicknamed number-crunching, warranted the purchase in 1972 of a medium-sized computer which, though about nine times slower than the 360/67 (now probably forty times slower than the University's current computer, the Amdahl 470), economically provides overnight service adequate for experimenters' purposes.

J. Wehrley Chapman of the Physics Department recently conducted a survey of peripheral computers around the campus. The goal of the survey, the latest in a series of surveys of this kind that have been made every two or three years, was to enable users of similar equipment to locate each other and trade insights into their machines. The list that Chapman has prepared makes it possible for prospective buyers of a given small computer to call on present users of that machine for advice about its merits and failings.

Keypunching Students

Many times above we have alluded to the use of the computer by students, who at Michigan can gain experience on a computer system that provides an unmatched variety of services and features. How did instructional computing begin?

When there were no central computing facilities during the fifties students were already using and even building computers at Michigan. From 1952 to 1957 the Electrical Engineering Department had its students building a small-scale electronic digital computer that was destined for some use and much educational tinkering. In the same period a young faculty member in mathematics, John Carr, who had a part in the development of the Willow Run Laboratory's MIDAC, was running student programs on that computer. Carr was very interested in using the new programmable computer for instruction and played the role of a courier, transporting batches of students' punched cards to MIDAC each week.Carr tried to persuade the staff at Willow Run to make its computer available to a wider number of students than those he could personally slip into the system, but regulations imposed on the use of the computer by federal project sponsors ruled this out. It was the tenacity of John Carr more than anything else that prompted the decision to acquire an IBM 650 for instructional use and to start the University's plans for an organized center of computing.

The fifties saw a relentless growth in the number of faculty and students who could program and use computers. The Ford Foundation Project on the Use of Computers in Engineering Education boosted this number greatly. An IBM survey in 1960 showed that the University was producing more computer-competent persons than any other university. By May, 1964, the *Research News* was able to report that the University had "perhaps the highest percentage of students and faculty members in the country who are able to prepare programs directly for the computer." It cited an estimate of 1,500 such persons on campus.

During the sixties the number of students using the computer and the number of courses requiring computer use grew enormously. This growth was stimulated by a grant from IBM to the University, over the years 1962 to 1967, for the purpose of "improving the lines of communication of the faculty and students of various disciplines with the automatic computer and of increasing their knowledge and productivity through its use."

In the first six months of 1960 sixty-three courses in sixteen departments required students to use the computer, while 114 doctoral students were using the computer for their dissertations. Eight years later, by 1968, this latter group had grown to 300 graduate students. Over three thousand other students were using the computer for course work in the fall of 1968, and 168 courses in thirty-three departments were dispatching their students to the Computing Center.

By the winter of 1970, 5,000 students in 266 courses and forty departments were on the computer. At this point, of course, computer-based courses were offered at all levels and in many departments. The object, naturally, of teaching the use of the computer was to enable students not just to solve exercises but to know how to apply the computer to new problems. In 1965 the National Science Foundation sponsored a U-M project to teach the use of computers in engineering design. Project directors Donald L. Katz and Brice Carnahan trained engineering teachers (both from the U-M faculty and from others) to exploit computers to solve design problems and to pass the techniques on to students who already had a background in computer-based courses.

Growth of computer use by students has not let up and does not promise to let up within the next few years. In the winter of 1974 almost 8,000 students enrolled in courses that would involve the computer. Between January and June of that year, 337 courses used the computer. The Department of Political Science had the largest number of computer-oriented courses with twenty-eight. During the same period there were nearly 350 dissertations for which data and programs were filed at the Computing Center.

To look at figures like these from the computer's viewpoint, the load of instructional computing amounts to about 100,000 jobs per month. Of course many of these jobs are false starts in which the computer spends a tiny fraction of a second detecting a programming error and counseling the student about what to do next. (In 1961 when the MAD language was current, students needed 3.8 trials before their programs ran through completion; this figure was arrived at to show that MAD was a simpler language for many academic purposes than FORTRAN, use of which by students called for more than five trials on the average.) MTS finds the large number of student jobs no challenge to handle or account for. A program devised by Charles Engle at the Computing Center does the bookkeeping for all the jobs at the end of each month and debits accordingly each teaching unit's allocation of computer time or research project account.

Computing Convenience

During the twenty-five years of development that now lie behind computing at Michigan, there have been some conflicting views over how best to proceed. Not all faculty members have been pleased with those institutional decisions that were not entirely advantageous to themselves. To be sure, there are other ways than the Michigan ways of providing a university with high-speed computation. One university long ago formed a company that sells computer services both to itself and to other customers in the city where it is located. Another university has no central computing facility; everyone there buys computing wherever they can or, when research grants permit it, they buy their own small or medium computers. Some schools pool different kinds of computing facilities and connect them by telephone lines; in such *networks* a user on one campus can go on line to facilities at any of several schools. Other universities that have central facilities have sometimes erred in locking themselves into equipment that grew outmoded or too small with unexpected rapidity.

Faculty members who have differed with computing policy at Michigan have sometimes later gained firsthand experience with the policies and equipment at other universities. One of the gratifications of of Robert Bartels, A. G. Norman, and Charles G. Overberger (Norman's successor, the present vice-president for research) is hearing from these persons. Computing systems as good as the one at Michigan and computing policies as accommodating of various needs are not numerous.

The main reason for this excellence is the breadth and convenience of Computing Center services. Constantly striving to meet the needs of as many kinds of users as possible, the Center has provided among its own software ancillary programs of many types. Few users have to seek elsewhere for some special service that the Computing Center staff never planned for. At Michigan, users can get immediate hands-on control of their own data because subsystems are already on file to make it possible.

An example of a subsystem of unique value is MIDAS. A creation of the Statistical Research Laboratory, MIDAS is a program for handling a large collection of interrelated statistics programs that understand the same commands. One of MIDAS's users calls it a pushbutton data-massager. It is a fast, easy, and inexpensive way of doing virtually all the statistical operations on a set of data that virtually any user could want. MIDAS was called up on MTS and used 110,000 times in 1974. Operable only on MTS, where it originated, MIDAS is said to be far better than comparable statistical software on other operating systems.

Users often want a special computer language that is appropriate to their material. MTS provides twenty-three language processors in common use, like those for FORTRAN (a general purpose scientific programming language), PL 1 (a general purpose data processing language), SIMSCRIPT (a language for simulations), BASIC (a language for novices), and processors for some languages in less common use, like SNOBOL (a manipulator of strings of characters), LISP (another string processor), GPSS (a general purpose language for simulation techniques), ALGOL (a general purpose scientific programming language), and APL (IBM's advanced programming language).

MTS's convenience is not matched by corresponding costliness. Its rates are comparable to what one might pay to a commercial vendor of time-shared computing. Yet such a vendor is likely to provide only a small portion of the services and software that are available from the Computing Center.

One of the prices that must be paid for systems convenience and flexibility is a measure of compromise in the efficiency of the processor's operation. For the mix of academic and research uses that the Computing Center serves, processor efficiency under MTS could hardly be greater. Yet conceived narrowly, MTS entails inherent operating inefficiencies. A processor devoted solely to information storage, for example, can be operated at maximum efficiency for that purpose alone. Yet such maximization means great inefficiency along other programming lines, in statistical operations, for example. Computer graphics has still other special requirements that do not accord with other kinds of processing. MTS strikes a compromise in processor efficiency to meet many programming needs and does so with the highest possible systems efficiency.

New Processor, New Manufacturer

Another form of convenience enjoyed by computer users at Michigan is the secure sense that MTS can go on meeting their changing needs. Users can rely on the Computing Center to assimilate new hardware when the old shows signs of obsolescence. In 1973 the two IBM 360/67 processors became inadequate to serve the growing number of computer users. Almost daily between 9 A.M. and 6 P.M. the processors were getting one-hundred percent use, a sure sign that users were demanding more computer power than the system could supply. Responsive daytime use of MTS was hard to get, and for many users this was a critical frustration. It was no use having patient health records on file, for example, or planning on a classroom demonstration at the terminal if an overloaded MTS would not answer the telephone.

Here Computing Center director Robert C. F. Bartels discusses a technical advance in computer gear with Professor A. Geoffrey Norman, who was vice-president for research during many of the years most critical in the development of Michigan's excellent computation facilities. In Bartels' hand is a multiple chip carrier within which are the circuits that compute. The round buttons visible are carefully devised cooling studs and fins that keep the temperature inside the computer within allowable limits. The chip carrier is a module within the University's current computer, the Amdahl 470V/6.

For some technical reasons the addition of a third 360/67 processor would have been more expensive than the added capacity would justify and would at most have increased MTS's capacity by fifty percent. It might thus have proved necessary again in two or three years to make further hardware changes. Major new burdens were at the time on MTS's horizon, for example, the large-scale use of computer-aided-instruction (CAI) in the School of Dentistry. CAI was scheduled in 1975 to bring a large number of new terminals on line and to increase greatly the number of channel addresses that MTS must accommodate.

The Computing Center staff proceeded in several steps to solve its hardware problem. In January, 1975, the University sold the 360/67 computer system and replaced it with an IBM 370/168 having approximately three times the job-processing capability of the 360/67. The 370/168 computer was to be rented temporarily until an evaluation of a forthcoming newer and faster computer could be made. The decision to acquire the Amdahl 470V/6, which had been in design and construction for several years by the newly formed Amdahl Corporation, was based on cost considerations and on the machine's ability to equal or outmatch the speed and reliability of the installed 370/168. The decision also depended on the assurance through formal agreement that maintenance and spare parts would be available during its presence at the University. During the summer of 1975 the 470 was tested against the 370/168 by alternating (without undue inconvenience to users) the control of peripheral devices between the two machines. By the end of September, 1975, the Amdahl had proved its worth and had supplanted the 370/168.

U-M computer users were completely unaffected by software difficulties in the transition to the Amdahl 470. The conversion from MTS on the 360/67 to MTS on the 370/168 affected some users, but very few. The Amdahl 470 is by design compatible with the 370/168; both can serve the same software and peripheral controls. For the user, only the speed of the 470 distinguishes the new machine from the slower 370/168. The main software change imposed by the 470 was relevant only to the Computing Center's staff. The 470's error recovery features (the way a machine notes that it has erred, re-tries the step that failed, then reports it) are considerably more definitive than those of the 370/168. MTS was modified to exploit these features, and is now better able to recover from errors and provide detailed records on machine performance.

For the computational services required by the University community, the Amdahl 470 has on the average at least four and one-half times the processing speed of the duplex 360/67 computer system. This increase in speed, together with the expanded channel capabilities of the Amdahl 470V/6, provided the confidence to purchase this machine. Bartels expects that with timely additions of high-speed memory the 470 will adequately serve the University's needs over the six-year period during which the purchase is financed. Because of its speed and its extended channel capabilities, the Amdahl 470 is perhaps a better buy for U-M purposes than versions of IBM's 370/168 now available. But, as with the 370/168 with which it competed, one disadvantage of the 470 when compared with the duplex 360/67 is that it has but a single processor. The 360/67 was able to provide continued, albeit degraded, service even with one processor down. This is no longer possible with its replacement. However, the 470 has an alternative advantage of being more easily diagnosed and fixed. Technology has advanced since the 360/67 era, and the frequency of component failures in proportion to number of transactions performed has declined, while the packaging of these components has eased and speeded the replacement of failed parts.

Toward the 1980's

Thus equipped with a powerful new processor, the Amdahl 470, the Computing Center will continue to provide the service that MTS users have come to rely upon. No policy changes are anticipated in the near future. But the world of computing is changing all the time. What will computing at Michigan be like ten years hence?

According to Charles G. Overberger, who took over the research vice-presidency when Dr. Norman retired from that position in 1973, one possible change of course for computing at Michigan might be heavier involvement with networking. A computer network links several computers, each with special strengths, via telecommunications. A

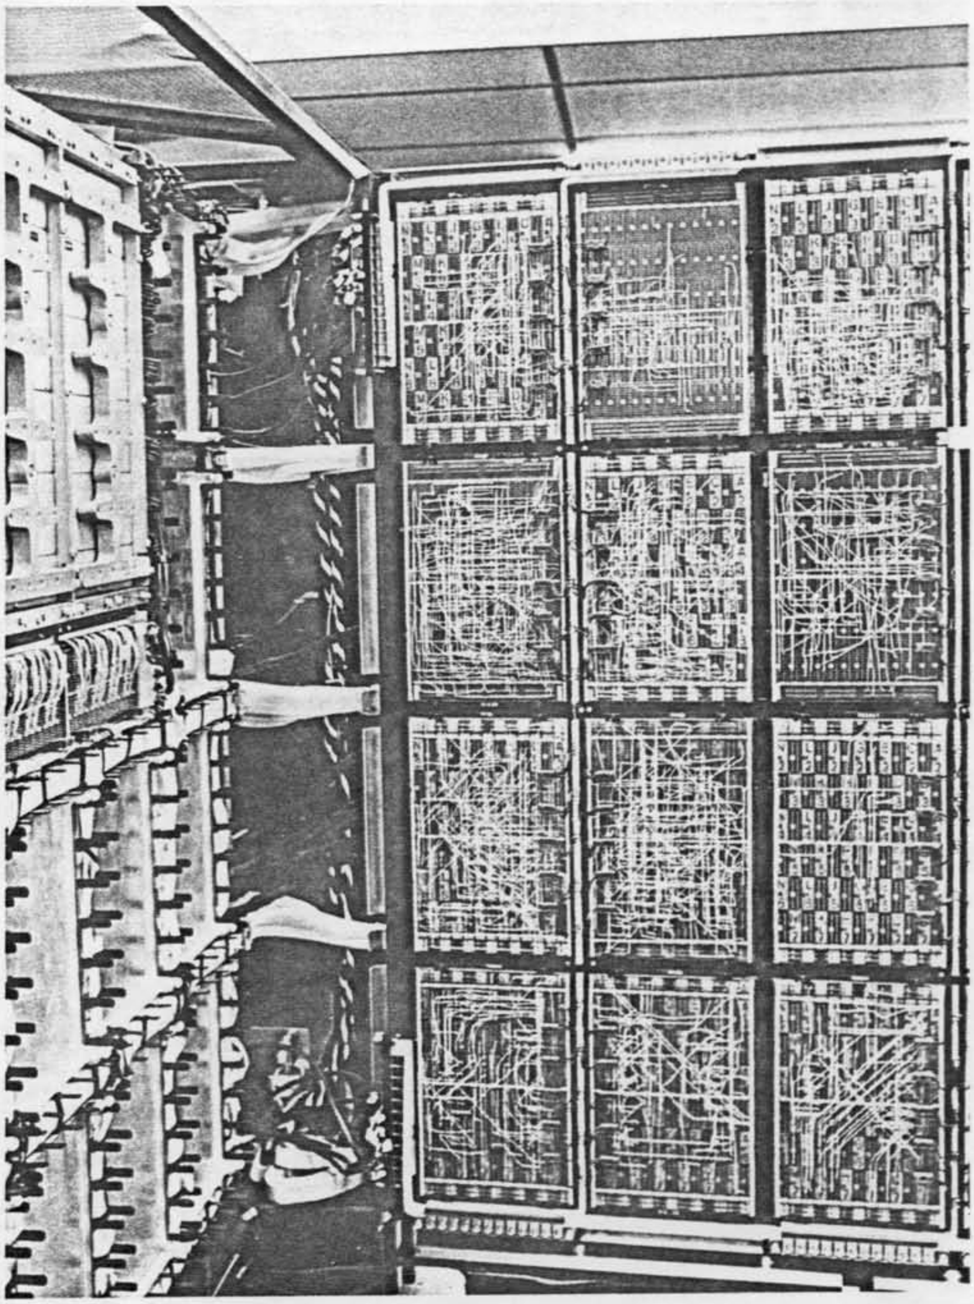

An inside look at a part of the University's newest central computer, the Amdahl 470V/6.

user at any of the network's units can contact the network computer that has the best software or best operating mode for his needs.

For some years Michigan has been involved with Wayne State University and Michigan State University in a network called MERIT. Planning for MERIT began in 1965, and in 1969 the Michigan Legislature and the National Science Foundation provided a development subsidy for the network. When MERIT began providing service in 1972, computer users at any of the three universities were able to reach another computer from their own office or laboratory terminal, and they could exploit any of that computer's special services. Like MTS, the system is fully interactive and convenient to use. A user need only type simple instructions into his terminal, and in a few moments his own local computer facility will have gotten him a line to a remote computer. The system also accommodates batch jobs, permitting their results to be printed out at any of the three sites on the system. Other colleges in Michigan, particularly Western Michigan University, Oakland University, and Grand Valley State University, have expressed interest in coming on line to MERIT.

One of the purposes of MERIT is to deliver to small institutions the computer sophistication that is already available at major institutions. This will help avoid duplication of facilities, particularly of software, which is ever becoming a proportionately more expensive aspect of computing. The three MERIT computers each have certain associated software not found at the other two. Michigan State, for example, has a geography research program called GPE, which deals with spatial data. It also has a compiler for the programming language PASCAL, which is especially suited for teaching people how to program and for constructing large programs. Wayne State's computer can be operated under OS as well as MTS and thus makes OS available to all MERIT users. Wayne State also has a special software system that allows for calculations out to 76,000 decimal places. And of course the University of Michigan (whose computer is the most frequently used in the network) also has a variety of unique software.

The operating modes of the three computers differ, and there are numerous instances in which a program will run appreciably cheaper on one than on the others. Eric Aupperle, MERIT's director, urges that researchers with repetitive programs to run try MERIT and perhaps find a computing bargain.

Networks are costly to develop, Overberger notes, and thus they need subsidies. But from a long range point of view they can lower costs by avoiding hardware and software duplication, as for example among colleges in one state.

But Overberger also points out that it would be unwise for an institution like Michigan to sacrifice wholly its substantial independence in computing. Today computing is at the heart of much research, and faculty members who are leaders in their fields must be assured that they will remain close to any decisions affecting their own computer needs. To proceed otherwise would jeopardize those benefits to society that result from research.

It will also be necessary to retain at Michigan a capacity for research in computing itself. This means that Michigan will continue to use the most advanced computing equipment and thereby prepare its students not for the computers of yesterday but for those of tomorrow.

Robert C. F. Bartels, Professor of mathematics and director of the Computing Center since its creation, talked with reporters when the U-M's most recent computer, the Amdahl 470V/6, was installed at the Computing Center.

Steps to Excellence

In reviewing the history of computing at Michigan we might ask, what were the key decisions that have today resulted in computing services and facilities that are so well suited to research and instructional computing? An early important decision was to separate academic computing from administrative bookkeeping and data processing; this led to research and instructional computing getting the priority it would need. Another important step was the creation of a central computing facility, which would serve all U-M users and thereby be a big enough system always to justify and pay for the best equipment. Later, when small computers specially suited to certain tasks became available, this policy was modified to allow their purchase, but only to the extent that a broad base of users for the central computer not be jeopardized. There was also wisdom in the decision, in the early years, to insist that computing and the Computing Center pay their own way; in this manner computing established a good reputation without putting University funds on the line.

A decision that has influenced all later decisions was to appoint Professor Bartels the director of the Computing Center; Bartels is a widely respected faculty member who has been able to balance all interests in the academic community and provide a widely useful computing system. The choice to pursue vigorously a time-sharing system gave a vital impetus to computing at Michigan, and though this decision asked at first more than technology could deliver, it ended in the creation of the Michigan Terminal System and a superb set of services tailored to the time-sharing mode and to the university environment. Another timely decision was to put University funds into the Computing Center's budget. Though in a way this decision was forced, it nevertheless accorded with the accepted policy that instructional computing and graduate student research computing should be fostered, not restricted, and with the belief that unit costs would decrease as volume of use increased.

The construction of a building specifically designed to house a computer continues to distinguish the University from most operators of computers, whose equipment is still housed in remodeled space. Finally the decision to replace the IBM 360/67 with equipment that would require few software changes has continued the U-M tradition of computing convenience. These decisions, coupled with perennially excellent staffing, have given Michigan a computing system second to none.

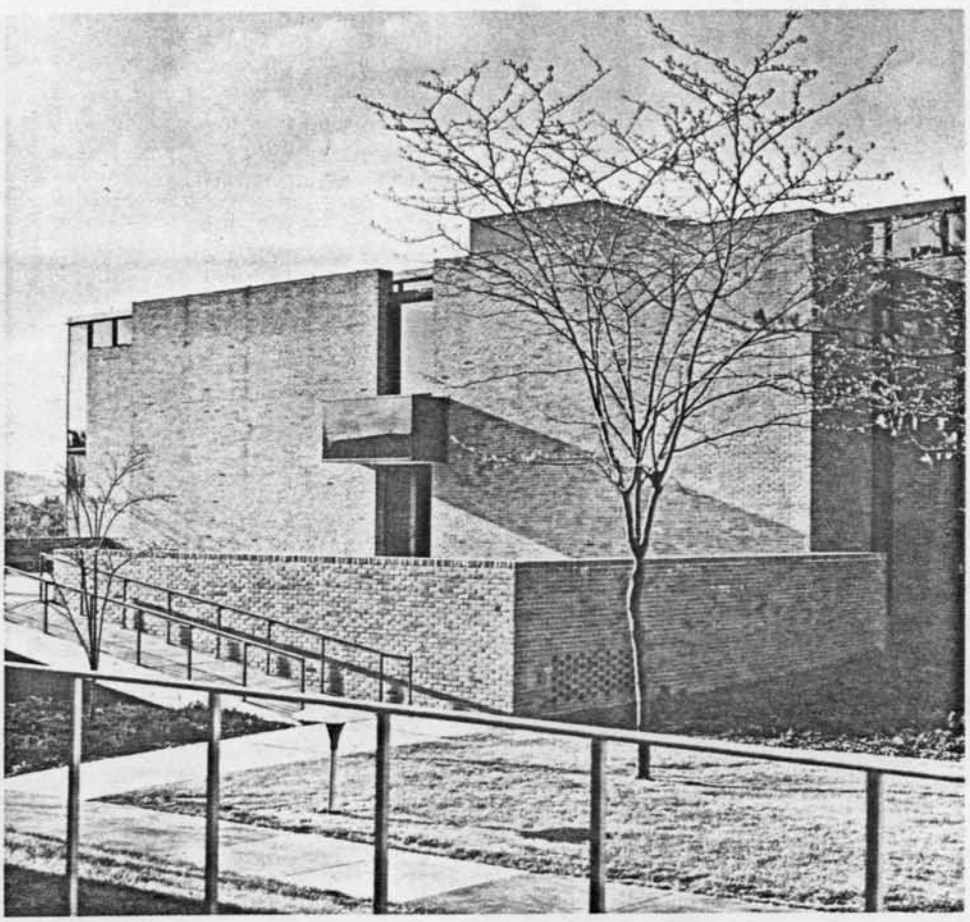

The north side and entrance to the Computing Center building is here captured evocatively in the dawn hours of a spring morning.

RESEARCH NEWS
Division of Research Development
and Administration
The University of Michigan
Ann Arbor, Michigan 48109

Address Correction Requested

Non-profit Organization
U.S. Postage
PAID
Ann Arbor, Mich.
Permit No. 144

The *Research News*, in describing a different area of scholarly activity each month, reflects the diversity of current research at The University of Michigan. The articles are designed to be interesting and understandable to the nonspecialist. Current issues of the *Research News*, including multiple copies for educational purposes, are available at no charge. Back issues dating from 1952 can be purchased from University Microfilms, 300 N. Zeeb Road, Ann Arbor, Michigan 48103. Correspondence concerning the *Research News* should be addressed to the Editorial Office, Division of Research Development and Administration, The University of Michigan, Ann Arbor, Michigan 48109.

Donald E. Thackrey
Editor and
Director of Program Development