
ENCYCLOPEDIA OF COMPUTER SCIENCE AND TECHNOLOGY

EXECUTIVE EDITORS

Jack Belzer Albert G. Holzman Allen Kent

UNIVERSITY OF PITTSBURGH
PITTSBURGH, PENNSYLVANIA

VOLUME 12

*Pattern Recognition
to Reliability of Computer Systems*

MARCEL DEKKER, INC. • NEW YORK and BASEL

COPYRIGHT © 1979 by MARCEL DEKKER, INC.

PROJECT MAC

Project MAC, renamed in 1975 the Laboratory for Computer Science, is an interdepartmental research laboratory in the computer sciences at the Massachusetts Institute of Technology in Cambridge, Massachusetts. It was organized in the spring of 1963 with the support of the Advanced Research Projects Agency of the Department of Defense under a research contract with the Office of Naval Research. Its founding director was Professor Robert M. Fano, who served in that position until September 1968. He was followed as director by Professor J. C. R. Licklider from 1968 to 1971, Professor Edward Fredkin from 1971 to 1974, and Professor Michael L. Dertouzos who has been serving since 1974. The acronym MAC was derived from two titles: Machine-Aided Cognition, expressing the broad objective of the research program, and Multiple-Access Computer, describing its major tool. The term "project" rather than "laboratory" was used to underline the interlaboratory as well as interdepartmental character of the research effort and, specifically, to encourage widespread participation on the part of individuals associated with other M. I. T. research laboratories.

The research and development program to be carried out by Project MAC was characterized in the first proposal as follows:

The broad, long-term objective of the program is the evolutionary development of a computer system easily and independently accessible to a large number of people and truly flexible and responsive to individual needs. An essential part of this objective is the development of improved input, output and display equipment, of programming aids, of public files and subroutines, and of the overall operational organization of the system. A second, concomitant objective is the fuller exploitation of computers as aids to research and education, through the promotion of closer man-machine interaction. The second objective is not only important by itself, but is also essential to the development of the computer system envisioned above, and vice versa. The third objective, which must be part of any university activity, is the long-range development of national man-power assets through education in the pertinent area: of faculty as well as of students, and outside M. I. T. as well as within the confines of the campus. Again, this third objective is inextricably interwoven with the preceding two, because people's approach to problems will have to evolve in parallel with the computer hardware and software.

The initial MAC computer system became operational in November 1963. It was a copy of the Compatible Time Sharing System [1], based on a modified IBM 7094 computer, which had been developed in the M. I. T. Computation Center under the leadership of Professor F. J. Corbato. The first version of this pioneering time-sharing system [2], on an IBM 708 computer, had been publicly demonstrated in November 1961 and a much improved version, including disk mass storage, was in daily operation on an IBM 7094 computer by the summer of 1963. The new MAC system was established to spur the evolution of an on-line research community, and it did not replace the similar computer installation in the M. I. T. Computation

[340]

PROJECT MAC

Center. The latter, which had to provide more traditional batch-processing service to the M. I. T. community, could be operated in a time-sharing mode only part of the time. The original MAC computer system was considerably improved and extended in the following 2 years and remained in operation until July 1973.

The goal of developing an on-line research community was achieved within the first 6 months of operation of the MAC computer system, as evidenced by the table of contents of the July 1964 progress report reproduced in Appendix A. By that time the system was serving some 200 users from 10 different academic departments. As many as 24 users could be using the system simultaneously from any one of some 100 teletypewriter terminals connected to the system through the M. I. T. telephone switching plant. Some of these terminals were located in the homes of researchers. Access to the MAC system could also be gained through the telex and TWX telegraph networks; tests and demonstrations had already been conducted from the West coast and even from European locations. The MAC system included, by then, the initial model of a multiple display system for computer-aided design [3] developed by the M. I. T. Electronics Systems Laboratory. This system included oscilloscope displays with character generator and light pen and also hardware facilities for rotating in three dimensions the image displayed. The MAC system itself was in continuous operation except for the periods of time required for preventive maintenance and for generating back-up copies on magnetic tape of the users' files stored on disks. Software facilities included some 60 system commands and amounted to some half a million words of machine code. Users' private files occupied, by then, most of two disk files with a total capacity of 18 million words.

The following quotations from a report on the MAC system [4], presented orally in June 1964, characterize some of the evidence that emerged from the first 6 months of operation.

Enthusiasm mixed with a great deal of frustration is the most common reaction to the system on the part of its users. The system was very quickly accepted as a daily working tool, particularly by computer specialists. This quick acceptance, however, was accompanied by the kind of impatience with failures and shortcomings that is characteristic of customers of a public utility.

It is becoming increasingly evident that the system's ability to provide the equivalent accessibility of a private computer is only a secondary, although necessary, characteristic. What users find most helpful is the fact the system places literally at their fingertips a great variety of services for writing, debugging, and compiling programs, and facilities for working on problems in their own particular field through the use of appropriate problem-oriented languages. The system users themselves are beginning to contribute to the system in a substantial way by publishing their work in the form of new commands. As a matter of fact, an editorial board is being established to review such work and formally approve its inclusion in the system. Thus, the system is beginning to become the repository of the procedural and data knowledge of the community that it serves. A corollary to this trend is, of course, the increasing difficulty that users find in ascertaining just what facilities of interest to them are included in the system.

The most delicate aspect of the operation of a multiple-access system of the MAC type is the responsibility assumed by the system managers with respect to the users' programs and data that are permanently stored in the

disk files. Elaborate precautions must be taken to protect the contents of the disk files against malfunctioning of the system as well as against actions of individual users. The supervisory program restricts the access of each user to his own private file and to public files, which he cannot modify. A full copy of the contents of the disk file is made twice a day to minimize the loss in case of malfunctioning. While losses of users' programs and data have occurred, the frequency and seriousness have not discouraged users from entrusting their work to the system.

The experience with the present MAC system suggests a trend toward memory-centered (as opposed to processor-centered) systems, including pools of bulk memories, core memories, central processor, and input-output channels, all communicating with one another, with core memories acting as buffers. On the software side, the trend seems to be in the direction of executing processes that consist of many subroutine and data structures, which are never assembled into a single program and some of which may be common to other independent processes simultaneously in execution. This view of computer systems is indeed very different from the traditional one. Its implications are far from clear. Their exploration is a major objective in the development of the next MAC system.

Planning on the follow-up MAC system began in earnest in the spring of 1964 with conferences with many computer manufacturers. It eventually led to the development of the Multics System in collaboration with the Bell Telephone Laboratories and the Computer Department of the General Electric Company, which later became part of Honeywell Information Systems, Inc. The details of this major research and development project are discussed below. Its goals and aspirations were presented in a series of papers at the 1965 Fall Joint Computer Conference [5], and the first version of the Multics System [6] became operational for general use at M. I. T. 4 years later in October 1969. The Multics System eventually became a commercial product of Honeywell Information Systems, Inc.

The first MAC computer system and the Multics development project were the most visible parts of Project MAC throughout the 1960s. However, Project MAC included from the very beginning a broad spectrum of research activities, as evidenced by the tables of contents of Project MAC Progress Reports in Appendices A, B, and C. The reader is referred to other articles and to the literature for information about other Project MAC research and for time-sharing research carried out elsewhere.

BACKGROUND AND CONTEXT

The initial success of Project MAC was largely a consequence of its very fertile environment. The goal of "machine-aided cognition" through the use of "multiple-access computers" represented the confluence of several lines of research and development in various M. I. T. laboratories, the most important of which are listed below.

1. The development of the SAGE Air Defense System at Lincoln Laboratory in the 1950s made available a variety of hardware and software techniques for on-line interaction with computers. More importantly, perhaps, it generated considerable

report stated further that "The conclusion that a central computer with remote consoles is required is based not only on the important economy obtained, both in running expenses and in programming expenses, but also on the importance of establishing close communication between the scientist and the machine." John McCarthy, who played the leading role in the preparation of the final report, added a further dimension to the notion of time-sharing in a lecture [15] delivered, also in the spring of 1961 as part of the celebration of the M. I. T. Centennial, by introducing the idea of computing as a public utility.

If computers of the kind I have advocated become the computers of the future, then computing may someday be organized as a public utility just as the telephone system is a public utility. We can envisage computing service companies whose subscribers are connected to them by telephone lines. Each subscriber needs to pay only for the capacity he actually uses, but he has access to all programming languages characteristic of a very large system. The system could develop commercially in fairly interesting ways. Certain subscribers might offer service to other subscribers. One example is weather prediction. . . . Other possible services include those specifically connected with computing such as programming services. Some subscribers perhaps might rent the use of their compilers. Other subscribers might furnish economic predictions. The computer utility could become the basis of a new and important industry.

Project MAC was organized and funded as part of a national program in information processing conceived by J. C. R. Licklider when he became director of the Information Processing Techniques Branch of the Advanced Research Project Agency of the Department of Defense. He was convinced, on the basis of his professional experience as a psychologist and computer user, that on-line interaction with computers (or, as he called it, "man-computer symbiosis") would lead to major steps forward in the utilization of computers in both civilian and military applications. The program was intended to be a broad but coherent research and development effort with long-term funding to be carried out in a number of university, nonprofit, and industrial laboratories. The time was ripe for the program, and Licklider's contagious enthusiasm provided the essential catalyst. Project MAC was conceived the day after Thanksgiving 1962, and its organization, objectives, and scope were informally agreed upon by the following Thursday. A formal proposal was submitted in January 1963, and a contract with an initial funding of \$2,220,000 was awarded by the Office of Naval Research on behalf of the Advanced Research Project Agency. Expenditure of these funds was authorized as of March 1, 1963. Project MAC's staff began moving into its quarters on the eighth and ninth floors of the just completed Technology Square Building in June 1963.

The first major activity of Project MAC was a 6-week summer study during July and August 1963. A total of some 57 people from universities, industry, and government participated in the study, with many of them in residence for the full 6-week period. The participants had access to terminals of the Compatible Time-Sharing System at the M. I. T. Computation Center, which by that time could support up to 21 simultaneous users. For many of them, this was their first opportunity to use a time-sharing system and to explore the potential of on-line, interactive computation. The reaction was very favorable in spite of the fact that the system was often overloaded and therefore slow in responding to service requests. The System

[344]

PROJECT MAC

Development Corporation, which was also conducting a large time-sharing project with ARPA's support, made its time-shared computer [16] (located in Santa Monica, California) available to the summer study participants via the Bell System TWX Network. The most valuable result of the summer study was the technical interaction that took place and the resulting working-level contacts that were established between many university, government, and industrial laboratories. The summer study became, in a sense, the launching pad for the overall ARPA program by establishing a community of people with a common interest in time-sharing and on-line interactive computation.

THE MAC COMPUTER SYSTEM

The initial MAC Computer System, a copy of the Compatible Time-Sharing System in existence at the M. I. T. Computation Center, became operational on October 1, 1963. It employed an IBM 7094 computer installation which included, in addition to the usual complement of peripheral equipment, two banks of 32,768 words of core memory, disk and drum storage units, a high-rate direct data connection for use with a graphical display system, and a transmission control unit for connecting the system to teletypewriters and other low-rate terminals through telephone lines. One of the two memory banks was dedicated exclusively to the time-sharing supervisor and other system programs that had to reside in core memory at all times, while the other bank was fully available to store user programs and data. The equipment included unique modifications to accommodate the two banks of core memory and a real-time clock, and to provide read and write protection through boundary registers to user programs while in core memory. The disk capacity grew during the following 2 years to 39 million words, and drum storage capacity to 559,200 words equally divided among three units. The drums were used to swap user programs in and out of core memory and to store directories and other frequently used data and programs. The graphical display equipment developed by the Electronic Systems Laboratory as part of the computer-aided design project became an integral part of the MAC system by the end of 1963. It included hardware for light-pen tracking, incremental digital line generation, and image scaling and rotation in three dimensions. A second independent display station, time-sharing the display generation hardware, was added a year later. By the middle of 1965, usage of each of the two stations averaged about 16 hr/day. Also, by the middle of 1965, the total number of teletypewriter terminals had grown to 157, all of which could access, through the M. I. T. PRX, either the MAC system or the similar computer installation at the M. I. T. Computation Center. The structure of the MAC computer installation as of the end of 1965 is illustrated in Fig. 1.

The time-sharing software of the MAC System was restructured and extended in a major way during the first 2 years of operation. A new file-management system was implemented, amounting to approximately 15,000 words of new code. It provided means for sharing private programs and data files among users without duplication and without danger of their being damaged. Permission by file owner could be selectively granted to (and withdrawn from) specific individuals, all members of a research group, or all users of the MAC system. The new file system also permitted the generation of back-up copies of any file or any modified file during normal operation of the system. It also provided on-line access to magnetic tapes and facilities for requesting, from teletypewriter terminals, the batch

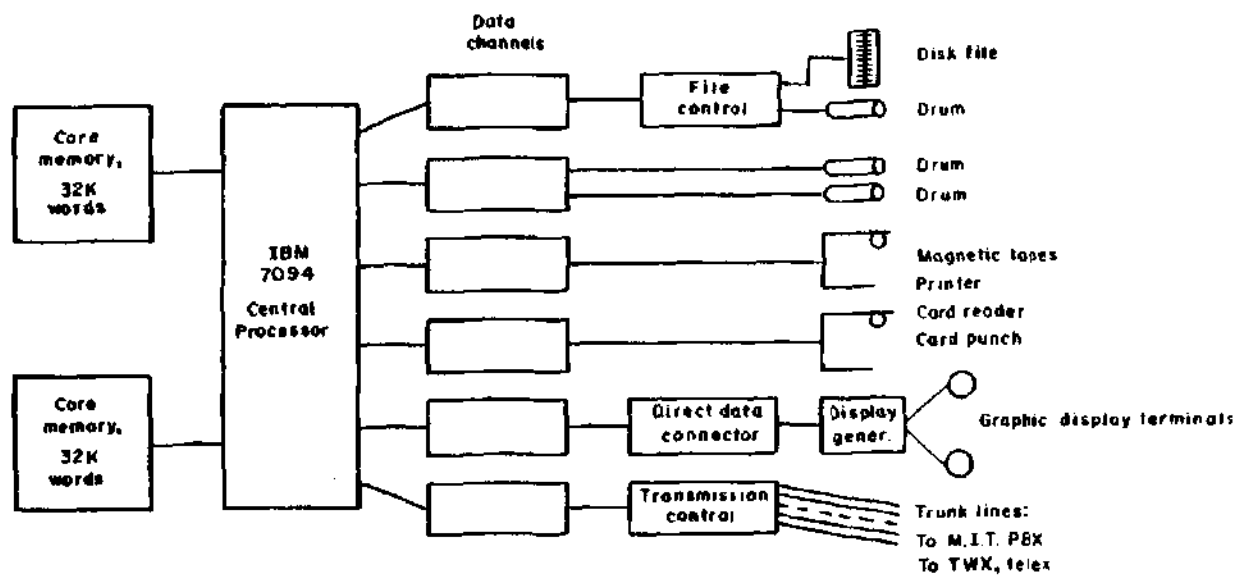


FIG. 1. Structure of MAC computer installation at the end of 1965.

[346]

PROJECT MAC

processing of programs in the background; that is, utilizing any processor time left over after meeting the requests of on-line users. The knowledge acquired in the development of this new file system was to pay significant dividends in the design and implementation of the Multics System. Another important innovation was a pair of commands, TYPSET and RUNOFF, for editing and formatting English text. These two commands became popular very quickly and were used for a variety of purposes including accounting and the maintenance of administrative data in Project MAC. The availability of these two commands made possible on-line storage and updating of system documentation, with users being able to print out for themselves updated copies of any section of the system's manual [17]. Two other very popular innovations were the command for linking teletypewriters to achieve intimate collaboration between users, and the mailbox feature for exchanging messages between users. Several new capabilities were implemented in the same period for facilitating the administration and monitoring of the system operation. The administration of the system was decentralized by allowing leaders of research groups to authorize new users, change passwords, and assign and reassign computer time and storage space to members of their groups within limits specified by the central administration and enforced by the system. Individual users could, in turn, check the state of their account at any time, and were automatically informed, upon completion of each command, about the computer time used by that particular command. Metering capabilities made it possible to tune the system to the point that 30 users could be allowed to work simultaneously. The scheduling algorithm was designed to make the total response time proportional to the actual computation time so that short requests were usually satisfied with a negligible wait time. A "party-line" scheme was employed to allocate access to the system among users when service was requested by more than 30 of them. These and other administrative features proved to be essential to proper management of the system because of the large, diverse, and physically dispersed user community, which included people located at universities on the West coast.

THE EARLY RESEARCH COMMUNITY

The objective of Project MAC, as stated in the original proposal, required the active participation and cooperation of a community of imaginative on-line computer users engaged in research, problem solving, and decision making in a wide variety of fields. Such a community of users was expected to contribute directly to the evolution and enrichment of the MAC system through the development of problem-oriented languages and other programming aids. This indeed happened, and very quickly.

The availability of the MAC Computer System and of support for on-line interactive research resulted in a sudden explosion of computer research on the M. I. T. campus. It is impossible to convey its scope in a few sentences, except by quoting statistics and mentioning a few examples. The second Progress Report of Project MAC covering the academic year 1964-1965, whose table of contents is reproduced as Appendix B, includes reports by 130 different members of the faculty, research staff, and students from 14 academic departments and 4 interdepartmental research laboratories. Some of the research topics reported on fall under the heading of computer sciences; others, while contributing in a substantive way to the goals of Project MAC, were primarily motivated by objectives outside the computer field.

The same progress report lists 18 MAC Technical Reports, 67 student theses, 47 journal and conference articles, and an equal number of other M.I.T. technical reports. On-line research with substantive external goals ranged from the development of problem-oriented languages in civil engineering to social systems analysis, from molecular model building to library information retrieval, from speech analysis to plasma physics, and from mathematical analysis to industrial dynamics. In the area of artificial intelligence, the foundations were being laid for what became, eventually, the MACSYMA system for symbolic mathematics, for the development of visually controlled manipulators, and for game-playing programs. In other computer science areas, work was in progress on the ELIZA system, a natural language, on-line, man-machine conversational system. An early version of this system was already in use at the Massachusetts General Hospital for psychiatric research. Research was also in progress on resource sharing in time-shared computer systems and on program segmentation, which was to play a major role in the design of the Multics system. The development of the AED (Automated Engineering Design) system, carried out in collaboration with the Electronic Systems Laboratory, was the major language development effort of the period. The AED language, an extension of ALGOL, was to become a major software development tool. Work had also begun, in the same period, on a storage tube display terminal which was eventually produced commercially.

The Project MAC research staff grew rapidly to 236 by July 1964, to 338 by July 1965, and peaked at 400 in 1966 and 1967. About two-thirds of the staff consisted of faculty and students. This community of on-line researchers contributed in a major way to the MAC system. By the end of 1966, more than half of the system commands were user developed. They included a variety of problem-oriented subsystems for simulation, mathematical analysis, text editing and formatting, and engineering design in various specialties (civil, mechanical, nuclear, electronic, etc.). Furthermore, the ease with which access to private files could be selectively granted to others (and withdrawn from them) led to an unprecedented utilization of other people's work. By the end of 1966, each system user had an average of 35 private files and 30 access links to other users' files.

THE MULTICS PROJECT

The original Project MAC research proposal stressed the inadequacy of the IBM 7094 computer installation as a basis for serious time-sharing system research. This machine, the last model of a venerable line of scientific computers, was never intended to support a time-sharing system. It had serious limitations, such as the very limited addressing capability, originating from technological restrictions and a corresponding view of computers that were no longer valid. As a matter of fact, it was very surprising and fortunate that sufficient equipment modifications could be made to support a time-sharing system with the capabilities of the final version of the Compatible Time-Sharing System. A successful time-sharing system could not have been developed without the addition of the second bank of core memory inaccessible to user programs, a feature that required significant modifications of several units of the computer installation. The 7094 computer did possess, however, two standard features that proved to be of critical importance: high reliability resulting from full transistorization, and the availability, beginning in 1963, of a

(348)

PROJECT MAC

movable-arm disk file. The availability of drums was also very important, but not as critical as that of the disk file.

The search for a more suitable computer installation began in the fall of 1963. Contacts were established with all manufacturers of large computers, several presentations were held at M. I. T., and visits by Project MAC staff were arranged to manufacturers who had expressed an interest in collaborating with Project MAC. It was made abundantly clear from the beginning that Project MAC was looking for more than just equipment; it was looking for a manufacturer sufficiently interested in time-sharing systems to collaborate with Project MAC in the development of significant equipment modifications and additions to meet Project MAC's needs. The requirements specified by Project MAC included:

1. Read and write protection of user programs
2. Privileged instructions inaccessible to user programs
3. Direct addressing of at least 250,000 words
4. A multiprocessing capability with all processors playing identical roles in the system
5. An effective telecommunication unit with interfaces to high-data-rate graphic display terminals as well as conventional telephone lines
6. Mass storage units including a fast drum for transferring programs in and out of core memory
7. Hardware for efficient program paging [18] and segmentation [19, 20], including a suitable content addressable memory to reduce fetching overhead

The last requirement could not be met by any commercial computer existing or contemplated at the time, and therefore would entail major equipment development. It was, however, regarded as essential by the technical committee appointed to select an appropriate computer installation for the future. Paging was regarded as essential to achieve effective memory management in a time-sharing environment, and segmentation (two-dimensional addressing) was regarded as equally essential to the implementation of sharing on the part of several users of single copies of procedures and data. These features were in turn regarded as necessary elements of an effective and efficient time-sharing system.

By the summer of 1964, specific proposals for computer installations were presented by three manufacturers: Digital Equipment Corporation, the Computer Department of the General Electric Company, and IBM Corporation. The technical committee evaluated the advantages and disadvantages of the three proposals and came to the unanimous conclusion that the proposal presented by the Computer Department of General Electric Company came closest to meeting the requirements of Project MAC, and a contract for the procurement of the computer installation was signed in August 1964. The installation was to be based on the GE 635 computer and was to include additions and modifications to meet the last of the requirements listed above. These additions and modifications, originally intended to be implemented as special features of the Project MAC installation, were eventually incorporated into a new computer, the GE 645. In addition, a fast, large-capacity drum and a general input-output controller were designed in collaboration with Project MAC and eventually became integral parts of the GE 645 computer system. The hardware specification phase was completed by fall of 1965.

In early 1965 the Bell Telephone Laboratories decided to acquire the same computer installation as Project MAC and to participate in the software development

effort. Shortly thereafter, the Computer Department of General Electric Company decided to join Project MAC and the Bell Telephone Laboratories as a third partner in what was to be known as the Multics Project. The partnership was an informal one, based primarily on the common technical objective of developing a time-sharing system capable of providing MULTiplexed Information and Computing Service to a large, diverse community. A series of papers [5] outlining the objectives and overall organization of the Multics System, authored by members of the three organizations, were presented at the 1965 Fall Joint Computer Conference.

The Multics Project turned out to be longer and much more difficult and demanding in terms of both manpower and resources than anticipated. It was, however, successfully completed without compromising, in any significant way, the original objectives. The Multics System became operational at M.I.T. in October 1969, at which time responsibility for its operation was transferred to the M.I.T. Information Processing Center. The Bell Telephone Laboratories terminated its participation in early 1969. However, collaboration between Project MAC and Honeywell Information Systems (which had acquired the Computer Department of General Electric Company) continued until 1977. This collaboration resulted in major improvements of the Multics System with respect to both hardware and software, particularly in the area of system security. The Multics System has been offered as a commercial product by Honeywell Information Systems since 1973.

By the end of 1971, 2 years after its becoming available at M.I.T. for general use, the Multics System was serving a user community of some 500 people. It was operating 24 hr a day, 7 days a week. Its full equipment configuration, including two processors and three memory modules amounting to a total of 384 kilowords, was able to serve some 55 simultaneous users with satisfactory response time. This full system, whose structure is illustrated in Fig. 2, was normally split during off-peak load periods into two separate systems, a service system and a development system. By that time the Multics software, including compilers, commands, and libraries, consisted of about 1500 modules amounting to a total of some 1,000,000 words of compiled procedure code. The development of this software had required an estimated 200 man years of design and implementation work. The Multics System Programmer's Manual amounted to some 3000 typewritten pages.

The Multics System and its development are well documented in the literature [6, 21]. However, the following features deserve special mention.

1. Multilevel Memory. The Multics System incorporated, from the start, a three-level, automatically managed memory consisting of main memory, magnetic drum storage (later replaced by bulk core memory), and disk file storage. However, the entire memory system appeared to a user as a single-level memory consisting of procedure and data segments. References could be made at any time to such segments as if they resided in main memory. This feature provided to the user convenient access to all procedures and data stored in the system, while still permitting efficient and economic utilization of a variety of storage devices.

2. Protection of Data and Procedures. Multics permitted controlled sharing of segments without duplication of the object shared. Access to a segment could be granted by its owner independently for reading, writing, and execution only, or for any combination of these three access modes. An additional control mechanism permitted the placing of a segment in any one of a number of concentric "rings of protection" which provided a generalization of the usual protection relationship between supervisory and user programs. References to a segment from some other segment residing in an outer ring were limited to specific entry points, as is the

[350]
PROJECT MAC

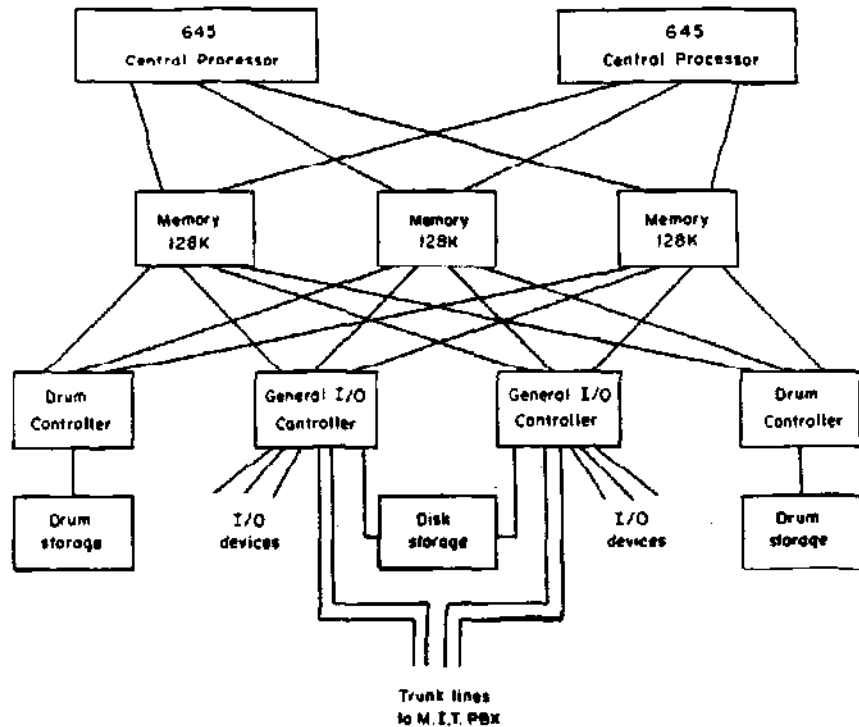


FIG. 2. Structure of the Multics computer at M.I.T. in 1971.

case with calls to supervisory programs on the part of user programs. The rings of protection mechanism was originally implemented in software, but a hardware implementation [22] was incorporated in the second-generation Multics processors that replaced the original one in 1972. Because of this feature, user groups could create, within Multics, closed subsystems to meet their own specialized needs with the protection attributes enjoyed by the Multics supervisor. For instance, user groups could be readily restricted to different programming environments suitable to their needs and capabilities. The rings of protection mechanism has special implications with respect to the protection of information privacy. It permits implementation of efficient data systems in which the utilization of data is controlled and monitored by procedures designed to fit the needs or privileges of different users.

3. On-Line Reconfiguration. Continuous operation was one of the original goals of the Multics System. To this end the system was designed to include at least two hardware units of each type, so that any unit could be taken out of service, because of preventive maintenance or failure, without interrupting the operation of the system. The original Multics operating system did not include the capability of reconfiguring the system without interrupting its operation. Facilities for dynamically adding and removing central processors and memory modules became available in 1971 [23], and from then on, dynamic reconfiguration was routinely employed at M.I.T.

to divide the hardware available into a service system and a totally separate development system during low user-demand periods.

4. Implementation in a Higher-Level Language. The Multics System was written in PL/I [24] except for some 17% of the modules, mostly databases or small subroutines executing privileged instructions, which were written in assembly language. The decision to use a higher-level language, specifically PL/I, was made at the very beginning of the Multics Project and turned out to be a very critical one. Implementation in assembly language would most likely have prevented the extensive redesign and rewriting of modules that proved to be essential to achieving the desired level of performance. On the other hand, the use of PL/I required the development of a series of compilers of increasing scope and performance which placed an additional major load on an already overburdened programming staff. The decision to program in PL/I continued to pay substantial dividends in terms of clarity and maintainability of the system long after its formal completion.

5. Design Iterations. The Multics Project, from inception to public operation of the system at M. I. T., including the equipment design and specification phase, took approximately 5 years, about twice as long as initially expected. The underestimation by a factor of 2, not uncommon in computer system development, was largely the result of not taking into account the time required for design iterations. In most engineering projects, design iterations take place primarily when detailed specifications are drawn, long before actual implementation. In software development, however, correspondingly detailed specifications require the use of a programming language. Therefore, the design and implementation phases have been traditionally confused, with the result that the need for design iterations is almost always overlooked. In fact, most modules of Multics were written two or three times and the design modifications were far from minor. It was only when major parts of the system were tested that unwise design choices became apparent; simulation would not have helped, as the knowledge necessary to construct an appropriate model of the system was not available at the time. The fact that design and implementation iterations are unavoidable when breaking new ground is one of the most important lessons resulting from the Multics Project.

6. On-Line Implementation. The entire development of Multics was carried out on-line. The early phase was carried out using the original MAC system, which not only facilitated coding but also provided a very effective communication mechanism among the design and programming staff of the three participating organizations. Much of the material pertinent to the project was kept on-line, readily accessible for reference and evaluation, and private messages and comments could readily be exchanged through the mail box mechanism. The MAC system was gradually replaced as a development tool by Multics itself as soon as it became usable by the system programming staff. The evolution of the system continued to be carried out on-line after it was opened to general use, and new modules were routinely installed without interrupting the operation of the system. Many key members of the Multics Project believe that the project could not have been carried to a successful conclusion if a general-purpose time-sharing system with the capabilities of the MAC system had not been available from the very beginning. Whether this was actually the case or not, there is no doubt that the Multics Project provided an excellent example of how intimate collaboration can be achieved through a time-sharing system.

(352)

PROJECT MAC

REORIENTATION OF PROJECT MAC

The responsibility for operation and maintenance of the Multics System was turned over to the M. I. T. Information Processing Center in October 1969. That date, 7 years after Project MAC was conceived, marks the end of its initial phase with time-sharing as its primary focus. By that time, on-line use of computers was widespread and time-sharing service was being offered commercially by several companies. Also, many time-sharing systems had been developed elsewhere to fit a variety of needs and which spanned a broad range of sizes and capabilities. The original goal of Project MAC had been largely reached, and time-sharing was no longer a priority research topic.

An important consequence of this change of interests and priorities in Project MAC was that its research program became more narrowly focused on computer science and engineering topics, with faculty and students being drawn largely from the Electrical Engineering and Computer Science Department. Thus the character of Project MAC changed rapidly from that of a research effort involving people from a large number of M. I. T. academic departments and laboratories to that of a discipline-oriented research laboratory of a traditional type. The name was correspondingly changed, first to MAC Laboratory and later, in 1975, to Laboratory for Computer Science.

The character of the research support changed equally rapidly in the early 1970s. While Project MAC was originally funded entirely by the Advanced Research Project Agency of the Department of Defense, the Progress Report for the academic year 1974-75 lists six sources of research support in addition to ARPA: Office of Naval Research; National Science Foundation; Air Force (Rome Air Development Center); Department of Health, Education and Welfare; IBM Corporation; and Honeywell Information Systems, Inc.

While time-sharing was at the center of the stage, other computer science research areas in Project MAC were growing in importance. The artificial intelligence group, in particular, had become a major part of Project MAC and was being funded by ARPA under a separate contract. It became an independent, interdepartmental laboratory in December 1970; this administrative split did not, however, *diminish the intellectual ties between that group and the rest of Project MAC.*

The rest of Project MAC was reorganized into several divisions and groups. The table of contents of the 1973-74 Progress Report in Appendix C provides a snapshot of the spectrum of research activities just before Project MAC became, in name as well as in fact, a permanent part of the academic structure of M. I. T. The list of Technical Reports of Project MAC published up to that time and the complete list of articles published by Project MAC staff are too extensive to be included here; they can be obtained from the Project MAC Progress Reports listed in Appendix D.

APPENDIX A: TABLE OF CONTENTS OF PROJECT MAC PROGRESS REPORT TO JULY 1964

ARTIFICIAL INTELLIGENCE

A Visually-Controlled Manipulator

Unrecognizable Sets of Numbers

The Algebraic Approach to Finite Automata

The Design of LISP 2

Man-Computer Symbiosis
Mathematical Assistant
Heuristics of Theorem Proving in Group Theory
Computer Experiments in Finite Algebra
A Chess-Playing Program
MATHLAB: On-Line Symbolic Computation
INTEGRATE: On-Line Indefinite Integration
BIOLOGY DEPARTMENT
Molecular Model Building
CIVIL ENGINEERING DEPARTMENT
Computer-Aided Teaching
The Structural Design Language
Dynamic Structural Analysis
Soil Engineering Problem-Oriented Language
I/O System Research
Bridge Design
Optimal Synthesis of Road Networks
COMPUTER COMMUNICATION STRUCTURES
Sequences and the Four-Color Problem
A Design Language for Digital Systems
Man-Machine Communication
Research on the Theory of Automata
Program Segmentation
Waveform Transformation and Graphical Display
A Table-Directed Translator
On-Line Braille
Analysis of Time-Shared Computer Systems
Queueing Models for File Memory Operation
Input/Output Subsystems
Non-Repeatability of Multi-Process Computations
Semantics for Multiprogrammed Computations
Optimal Allocation of System Resources
Automatic Flowcharting
Visual Information Processing
COMPUTER OPERATION
CTSS Operation and Equipment
COMPUTER SYSTEM RESEARCH
Research on the Compatible Time-Sharing System
MULTICS Hardware System Design
ELECTRONIC SYSTEM LABORATORY
Introduction
Display Systems Research
Computer-Aided Design
Computer-Aided Electronic Circuit Design
Accelerometer System Studies
Simulation of Strapped-Down Navigation Systems
LIBRARY RESEARCH
Technical Information Project (TIP)
Process Control: Serials and Journals
Search Procedures

[354]

PROJECT MAC

Measures of Relatedness
Statistics of Words in Titles
Educational Use of TIP
Use of TIP to Update a Data Compilation

LINCOLN LABORATORY

Baseband Design of a Unified Carrier
On-Line Data Storage and Retrieval
Compilation for a Digital Differential Analyzer
On-Line Experimentation

NON-MIT USERS

Model Testing
Discovery and Learning Techniques and Programmed Instruction
Generalized Desk Calculator

OPERATIONS RESEARCH CENTER

Aspects of Integer Linear Programming

RESEARCH LABORATORY OF ELECTRONICS

Introduction
Dynamics of Active Plasma Systems
Dispersion Relation for Hot Plasmas
Plasma Dispersion Relations with Infinite Roots
Stability Analysis of Dispersion Relations
Use of CTSS in Plasma Physics Experiment
Analysis of Speech
Simulation of the Human Larynx
Articulatory Events of Speech Generation
Grapheme-to-Phoneme Translation of English
Programming Support and Development
RLS Statistics on CTSS
Sorting of Personnel Records

COMIT

SCHOOL OF ENGINEERING

Computer-Aided Design
On-Line Mathematical Analysis
Derivation of Preliminary Ship Lines
Time-Sharing Reactor Code System
A Stress-Analysis Program
Stress-Analysis Conformal Mapping
Computer-Aided Teaching of Dynamic Systems Behavior
Automatic Network Synthesis
A Computer Model of Kinesthesia
Computer Solutions for Boundary-Value Problems
Algebraic Expression Compiler
An Algorithm to Aid Logic Design

APPENDIX B: TABLE OF CONTENTS OF PROJECT MAC PROGRESS REPORT II,
JULY 1964 TO JULY 1965

ADMINISTRATION AND SERVICES

Measurements on the Time-Sharing System Performance

SCHOOL OF ENGINEERING

Educational Applications of Multiple-Access Computers
General Method for Network Synthesis
Command Programs for Time-Sharing
Computer-Aided Ship Design
Computer Evaluation of Cauchy-Type Integrals

CIVIL ENGINEERING DEPARTMENT

A Preliminary Study of On-Line Structural Design
STRESS: Structural Engineering Systems Solver
Transportation and Highway Systems
The Transportation Demand Project
Highway Route Location
Traffic Flow Analysis
Optimum Allocation of Traffic Flow
Traffic Simulation Studies
Railway Engineering Systems
Soil Engineering Problem-Oriented Language
Teaching Machine System
The COGO Language
Programming System for Project Scheduling
Unsteady Flow in Open Channels

RESEARCH LABORATORY OF ELECTRONICS

Dynamics of Beam Plasma Systems
Numerical Solution of Space Charge Wave-Propagation Constants
A Computer Display for Wave-Type Instabilities
Speech Analysis

SCHOOL OF HUMANITIES AND SOCIAL SCIENCE

Social Systems Analysis

SLOAN SCHOOL OF MANAGEMENT

A General Theory of Human Problem Solving
Time-Sharing in Psychological Research
Management Applications for the Development of Information utilities
Marketing Model Construction
Computer Evaluation of a Technique of Stock Trading
Computer Job-Shop Simulation

SCHOOL OF SCIENCE

Vortex Studies
Programs for Physical Problems
Energy Levels of Fluorine-77
Use of CTSS in a Plasma Physics Experiment

COMPUTER SYSTEMS RESEARCH

System Requirements for Time-Sharing
Research on the 7094 Time-Sharing System
Research on Systems Programming Languages

COMPUTER COMMUNICATION STRUCTURES

Computation Systems Development
Research of Scheduling Algorithms
Machine Structures
Grapheme-to-Phoneme Translation of English
Simulation of Large Computer Systems

[356]

PROJECT MAC

Simulation of Multiple-Access Computers
Measurement of Multidimensional Transducers
TREET: A New List-Processing Language
S-PLANE: A Program for the Manipulation of Zero-Pole Patterns

ARTIFICIAL INTELLIGENCE

Derivator
Computer Problem Solving with Natural Language Input
Mechanization of Proofs in Number Theory
The Mathematical Assistant
Integral Table Look-Up Procedures
A Geometry Theorem-Proving Program
File Maintenance and Syntax Generalization
Theorem Proving on a Time-Shared Computer
A Language for Binary Relations
M-Expression Translator

LIBRARY RESEARCH

Technical Information Project
Search Procedures and Measures of Relatedness

ELECTRONIC SYSTEMS LABORATORY

Display Systems Research
Computer-Aided Design Project
Input/Output Equipment for the PDP-1

LINCOLN LABORATORY

A Computer Technique for Optimization
A Method for On-Line Manipulation of Data Files

APPENDICES

A - Publications Other Than MAC Technical Reports
B - Project MAC Memoranda

APPENDIX C: TABLE OF CONTENTS OF PROJECT MAC PROGRESS REPORT XI,
JULY 1973 TO JULY 1974

FUNDAMENTAL STUDIES DIVISION

Introduction
Theory of Computation Group
A. Introduction
B. Inherently Difficult Problems
C. Algorithms
D. Foundations of Concrete Complexity Theory
Computation Structures Group
A. Introduction
B. Structured Programming Language
C. Proofs of Correctness
D. Algebraic Specification of Data Types
E. A Portable Compiler
F. Base Language Semantics
G. Data Flow Model
H. The Binding Model
I. Semantics of Data Structure

- J. Program Schemas
- K. Theory of Petri Nets
- L. A Universal Asynchronous Logic Array
- M. Data Flow Computer Architecture

AUTOMATIC PROGRAMMING DIVISION

Introduction

Automatic Programming Group

- A. Introduction
- B. Understanding How a User Might Interact with a Knowledge-Based Application System
- C. Attempting to Set Down the Knowledge Possessed by Expert Consultants
- D. Design of the OWL System
- E. Study of System Modeling an Analysis
- F. Study of the Process of Algorithm Generation
- G. Development of a System for Translating From a Very High Level Language into IBM/370 PL/I
- H. A Business Model for Automatic Programming
- I. Conclusion

Engineering Robotics Group

- A. Introduction
- B. Engineering Robotics
- C. Graphics
- D. Development of Timesharing System
- E. Language Semantics

Mathlab Group

- A. Introduction
- B. Description of MACSYMA
 - 1. Language and Interactive Facilities
 - 2. General Representations
 - 3a. The Rational Function Representation
 - 3b. Major Algorithms for Polynomials and Rational Functions
 - 4. The Integration Subsystem
 - 5. The Power Series Subsystem
 - 6. Miscellaneous Facilities
 - 7. The MACLISP System
- C. Summary

Clinical Decision-Making Group

- A. Introduction
- B. Funding
- C. Research in the Past Year
- D. Present Illness -- Protocol Collection and Analysis
- E. Present Illness -- Computer Simulation
- F. The Time Specialist
- G. Differential Diagnosis
- H. Digitalis Therapy Advisor

COMPUTER SYSTEMS RESEARCH DIVISION

- A. Introduction
- B. Certification of Computer Systems
- C. ARPA Network Activities
- D. Technology Transfer

[358]

PROJECT MAC

E. Other Activities

PROGRAMMING TECHNOLOGY DIVISION

A. Introduction

B. Programming Technology

1. CALICO

2. MUDDLE

3. Automation of Program Documentation

4. Project Reporting System

5. Application Programs

C. Networks

D. Other Activities

OTHER RESEARCH

PLANNER Group

A. Introduction

B. Procedural Embedding of Specifications

C. Theoretical Basis

D. Contracts

E. The Actor Model of Computation

F. A Brief Introduction to Plasma Syntax

G. Contract Examples

H. Intentions

I. Types and Constraints

J. Post-Requisites

K. Post-Conditions

L. Use of Intentions

M. Modular Distribution of Knowledge

N. Simplification Contracts

O. Abstract Operations of Queues

P. An Implementation of "Pure" Queues

Q. Program Verification Justification and Meta-Evaluation

R. Meta-Evaluation of the Implementation of Queues

S. Bug Detection Using Meta-Evaluation

T. An Alternative Implementation of "Pure" Queues

U. Overview of Meta-Evaluation

V. Benefits of Meta-Evaluation

W. Programming Style and Responsibility

X. Relationship to Automatic Debugging

Y. Relationship to Automatic Programming

Z. Relationship to Work on Programming Languages

A1. Relationship to Activation Record Semantic Models

B1. Relationship to Research on Theorem Proving

C1. Advantages of Contracts

D1. Overview of the Programming Apprentice

E1. Further Work

F1. Conclusion

G1. Acknowledgements

APPENDIX D: PROJECT MAC PROGRESS REPORTS PUBLISHED FROM BEFORE
JULY 1964 TO JULY 1973

Copies of all Project MAC reports listed may be secured from the National Technical Information Service, Operations Division, Springfield, Virginia, 22151. Prices vary. The AD number must be supplied with the request.

	<u>Publications</u>
Project MAC Progress Report I to July 1964	AD 465-088
Project MAC Progress Report II July 1964-July 1965	AD 629-494
Project MAC Progress Report III July 1965-July 1966	AD 648-346
Project MAC Progress Report IV July 1966-July 1967	AD 681-342
Project MAC Progress Report V July 1967-July 1968	AD 687-770
Project MAC Progress Report VI July 1968-July 1969	AD 705-434
Project MAC Progress Report VII July 1969-July 1970	AD 732-767
Project MAC Progress Report VIII July 1970-July 1971	AD 735-148
Project MAC Progress Report IX July 1971-July 1972	AD 756-689
Project MAC Progress Report X July 1972-July 1973	AD 771-428

REFERENCES

1. M.I.T. Computation Center, The Compatible Time Sharing System, M.I.T. Press, Cambridge, Massachusetts, 1963.
2. F. J. Corbato, M. M. Daggett, and R. C. Daley, An experimental time-sharing system, in AFIPS Conference Proceedings, Vol. 21, 1962, pp. 335-344.
3. R. Stotz, Man-machine console facilities for computer-aided design, in AFIPS Conference Proceedings, Vol. 23, 1963, pp. 323-328.
4. R. M. Fano, The MAC system: The computer utility approach, IEEE Spectrum 2, 56-64 (January 1965).
5. F. J. Corbato et al., Session 6: A new remote access man-machine system, in AFIPS Conference Proceedings, Vol. 27, Part 1, 1965, pp. 185-247.
6. E. I. Organick, The Multics System: An Examination of Its Structure, M.I.T. Press, Cambridge, Massachusetts, 1972.
7. W. A. Clark, J. H. Frankovick, H. P. Peterson, J. W. Forgie, R. L. Best, and K. H. Olsen, The Lincoln TK-2 computer, in Proceedings of the Western Joint Computer Conference, February 1967, pp. 143-145.
8. I. E. Sutherland, Sketchpad: A man-machine graphical communication system, in AFIPS Conference Proceedings, Vol. 23, 1963, pp. 328-346.

9. J. C. R. Licklider, Man-computer symbiosis, IRE Trans. Hum. Factors Electron. 1, 4-11 (March 1960).
10. W. Pease, An automatic machine tool, Sci. Am. 187, 101-115 (September 1962).
11. D. T. Ross, Automatically programmed tool system, Mech. Eng. 81, 59-60 (May 1959).
12. S. A. Coons, An outline of the requirements for a computer-aided design system, in AFIPS Conference Proceedings, Vol. 23, 1963, pp. 299-304.
13. D. T. Ross and J. E. Rodriguez, Theoretical foundations for the computer-aided design system, in AFIPS Conference Proceedings, Vol. 23, 1962, pp. 305-322.
14. J. B. Dennis, A multiuser computation facility for education and research, Commun. ACM 7, 521-529 (September 1964).
15. J. McCarthy, Time-sharing computer systems, in Management and the Computer of the Future (M. Greenberger, ed.), M.I.T. Press, Cambridge, Massachusetts, 1962, pp. 221-236.
16. J. Schwartz, A general-purpose time-sharing system, in AFIPS Conference Proceedings, Vol. 25, 1964, pp. 397-411.
17. P. A. Crisman (ed.), The Compatible Time-Sharing System: A Programmer's Guide, M.I.T. Press, Cambridge, Massachusetts, 1965.
18. T. Kilburn et al., One-level storage system, IRE Trans. Electron. Comput. 11, 223-235 (April 1962).
19. A. W. Holt, Program organization and record keeping for dynamic storage allocation, Commun. ACM 4, 422-431 (1961).
20. J. B. Dennis, Segmentation and the design of multiprogrammed computer systems, IEEE Int. Conv. Rec. 3, 214-225 (1965).
21. F. J. Corbato, J. H. Saltzer, and C. T. Clingen, Multics—The first seven years, in AFIPS Conference Proceedings, Vol. 40, 1972, pp. 571-583.
22. M. D. Schroeder and J. H. Saltzer, A hardware architecture for implementing protection rings, in ACM Third Symposium on Operating Systems Principles, October 1971, pp. 42-64.
23. R. R. Schell, Dynamic Reconfiguration in A Modular Computer System, Project MAC TR-86, 1971.
24. F. J. Corbato, PL/I as a tool for system programming, Datamation, pp. 68-76 (May 6, 1969).

Robert M. Fano

PUNCHED CARD DATA PROCESSING

INTRODUCTION

This article describes electromechanical computing devices originally designed for application to statistical and business accounting problems, but quickly adapted