THE PLACE OF TIME SHARING*

Robert M. Fano
Massachusetts Institute of Technology

Introduction

The experimentation with time-sharing systems during the last few years has generated considerable interest in their potential for research and education. Several articles have appeared arguing the pros and cons of these systems when compared with conventional batch-processing systems, and public debates have been staged at technical meetings. Unfortunately, the discusssion has been clouded by differences in meanings of terms, such as time sharing, and by differences in the specific systems, applications, and evaluation criteria that the discussants had in mind, but seldom made explicit. The spectra of time-sharing systems, batch-processing systems, applications, and evaluation criteria, are so broad that almost any statement about them can be true in one particular case and false in another. A recent article (1), which reports on an experimental comparison of on-line and off-line programming, makes abundantly clear how difficult it is to make meaningful objective comparisons even in a very restricted spectrum of situations. In any case, it seems to me that the dichotomy of time sharing versus batch processing has been unduly emphasized. In the first place, time-sharing service and batch-processing service can be provided simultaneously by the same system. In the second place, the main value of time-sharing systems stems from uses that would not otherwise be possible or practical. In the third place, the characteristics of computer systems, their overall economics and the practices in allocating costs are all changing very rapidly at this time. Thus, the proper technical issue on which to focus our attention is how to meet most effectively and economically the broad spectrum

*Work reported herein was supported by Project MAC, an M.I.T. research program sponsored by the Advanced Research Projects Agency, Department of Defense, under Office of Naval Research Contract number Nonr-4102(01).

of computer needs, fully realizing that a single system may not be adequate for all needs. On the economic side, the main issue is whether or not the benefits resulting from applications made possible by time-sharing systems warrant the cost of the necessary facilities.

This article will attempt to throw some light on the future role of time-sharing systems in engineering education. We will first discuss the nature of some of the frustrations currently experienced by both students and teachers, and how time-sharing systems may help in alleviating them and in raising the overall quality of education. We shall then survey the features of time-sharing systems which appear to be most important from a user's standpoint. Finally, we shall analyze the distribution of costs in the operation of a contemporary time-sharing system and make some extrapolations about the future feasibility of wide-spread use of time-sharing systems in engineering education.

The Role of Time Sharing

The role of time-sharing systems in engineering education is, in broad terms, essentially the same as in any other area of human activity. It is to bring the power of computers to the aid of each individual in whatever intellectual activity he may be engaged, by providing to him convenient access to a computer capable of acting as a knowledgable and skillful assistant. A significant part of this role hinges on the system's ability to communicate with many people at the same time, as well as at different times, and to facilitate intellectual communication between them.

Students as well as teachers and administrators are currently being frustrated by problems which are becoming increasingly severe and for which no solutions seem to exist within the framework of current educational practices.

Computers may provide a way, and possibly the only way, of breaking out of the constraints of the past, but only if they are made easily accessible to each individual, both physically and intellectually, in a manner intended to increase substantially his own intellectual capabilities. Large, general-purpose time-sharing systems are essential in this respect. Private computers are inadequate because, by definition, they can not facilitate communication between people.

Most of the problems that we are currently facing in engineering education, are not new. However, they are becoming increasingly serious as the number of students and the knowledge to which they must be exposed keep growing. The most serious problem stems from our inability to provide sufficient personal guidance to each student and to fit the educational process to his special capabilities and learning pattern. We all know that even in a selected class, in a highly-selective school, individual differences are quite evident, and that it is impossible to adjust the pace or the content of instruction to fit sumultaneously all students. Some students prefer to go from the specific to the general, and some from the general to the specific; some students must plough through the details before they can appreciate the broad picture, others can not force themselves to pay attention to details before they understand "what it is all about"; some students learn best by first skimming over the surface and then gaining depth by going over the same topic several times; others learn best by probing deep each point right away. Also, of course, some grasp new ideas quicker than others, some are motivated by different objectives than others, and some are just plain smarter than others.

A major objective of computer-aided instruction is, of course, to provide personalized instruction so as to permit each student to proceed at his own pace. Since computer-aided instruction is the main topic of a separate article in this issue, we shall consider only certain aspects of it which seem to depend critically on the availability of general-purpose time-sharing systems. Computer-aided instruction should involve much more than interaction between student and computer. It should involve interaction between a student, one or more instructors, and perhaps other fellow students, through a computer serving as a buffer and controller of the interaction. It is the responsibility of the instructor to select the paths through the subject matter to be followed by the students, to monitor their progress, to diagnose their individual difficulties, and to modify accordingly the rules within the computer program that control their work. In other words, an instructor should be able to guide, with the aid of the computer, the progress of each individual student, and to adapt the program with which the student is interacting to his individual needs. Furthermore, we all know that students learn a great deal from one another and that group activities are an important part of education. Thus, programs for computer-aided instruction should make appropriate provisions for group interaction. Also, a student should be able to get the attention of his instructor whenever the instruction program is unable to help him, and the instructor, in turn, should be in a position to intervene effectively from a remote location.

The point that needs stressing is that the computer system must enable the instructor to monitor the work of his students and intervene either immediately or at some later time. Also, the computer system must be able to identify each student, remember his past work, and control his present work according to the latest rules specified for him by the instructor. No existing time-sharing system is fully adequate in this respect. Another point worth stressing is that the rules embodied in current instructional programs for guiding the student's work are rather simple and are based mainly on the last response of the student. This implies that, while the student may be able to proceed at his own pace, no significant adaptation of the program to his learning pattern can take place without the specific intervention of the instructor. However, there are a few examples (2) of programs which can, through interaction with a person, build models of the situations that the person is attempting to describe. Thus, one can envision for the future instructional programs able to build models of the state of understanding of a student, that is, of what the student believes to be true. It may then be possible for an instructional program to diagnose most of a student's difficulties and guide his work accordingly without direct intervention of the instructor. Progress along these lines will probably be the limiting factor in the development of computer-aided instruction at the university level.

One last aspect of computer-aided instruction must be considered because of the requirements it places on the underlying computer system. We all know that knowledge can not be neatly packaged into separate boxes. Thus, the instructional program in a particular field will have to turn control over to programs in other fields just as a textbook in one field has to make references to textbooks in other fields. The need for turning control over to another program may come when a weakness in a student's background becomes apparent, or when a student wishes to pursue further a line of thought of particular interest to him. Thus, instructional programs can not be viewed as separate entities, but rather as forming a complex in which each program interacts in many ways with other programs. It should be possible, nevertheless, to modify each individual program without affecting the operation of the others. This implies that a program must be able to address other programs without knowing their physical locations, in a manner analogous to the way a textbook makes references to other textbooks. As we shall see later, the technique of "program segmentation" used in some of the newer computers designed specifically for time sharing, is intended to meet requirements of this type.

Another problem responsible for many of our present frustrations is that caused by ever-expanding knowledge. The time that a student can afford to devote to his education is finite while the knowledge that, at least in principle, he should acquire keeps growing. The crux of the dilemma is that we have no way of storing human knowledge in a directly-usable manner, regardless of how well established it is. When knowledge is stored in a book, a person has to read the book, absorb its content, and practice the pertinent intellectual skills before he can use it in any competent manner. Thus, for instance, we keep asking each generation of engineering students to practice symbolic differentiation and integration and to learn well-established techniques for solving differential equations in spite of the fact that this work adds very little to their understanding of the concepts of calculus. As a matter of fact, the pressure to learn and practice techniques is so great that many students become very skilled at integrating without ever understanding the concept of integral. Mathematical tasks such as symbolic integration can now be executed quite adequately by computer programs (3) without human intervention. These programs use essentially the same techniques that are taught to students; they inspect the integrand and make hypotheses about appropriate integration techniques, change approach when the integrand does not become simpler, perform substitutions of variables when appropriate,etc. We have here an example of how knowledge may be stored in a computer program in a directly-usable form. The important point is that computer programs may provide an alternate way of recording knowledge, with the advantage over books that a person does not have to learn and practice the knowledge so recorded in order to use it effectively. Of course, such programs can just as easily teach the techniques that they employ, if the person should be interested in learning them. This is only a beginning, but it suggests that future generations may be spared the effort of learning (i.e. rediscovering under guidance) knowledge developed by preceding generations. Of course, knowledge stored in computer programs will have then to be made very easily accessible to each individual. Thus, if this view of the future turns out to be correct, time-sharing systems will have to play a very central role in universities, much more crucial than the role currently played by libraries.

Let us turn our attention now to problems of university administration, which are becoming more acute as the size of the student population grows. By way of introduction, I would like to mention an experimental program (4) demonstrated some time ago to me at the System Development Corporation. This program was intended for on -line use on the part of high school students to help them in their academic planning, as a counselor would. The program seemed to be serving, quite effectively, two different functions. In the first place it was interpreting school programs, graduation requirements, and academic regulations for the benefit of each individual student, that is it was providing a sort of personalized catalogue of the school. In the second place, it was providing each student with statistical information about the college performance of students generally similar to him. The program seemed to perform these two functions rather well, and few advisers could have provided from the top of their heads as much detailed and pertinent information.

It is important to note that such programs while helping students, could obtain from them information about their plans, which could then be used by school administrators to prepare, for instance, instructor and class assignments. We can see here the beginning of an on-line school information system in which each individual is helped by programs and data stored in the system; in the process of being helped, each person provides information about his activities and plans which is, in turn, useful to other people. The important point is that information is provided by each individual in the course of being helped rather than as a separate action. Here again we see a computer system emerging as an assistant to the individual and as a powerful tool for facilitating intellectual communication between people. It is worth pointing out in this respect that the practical problems of keeping track of and planning for a large number of students is influencing and constraining in many ways our academic policies. An information-rich community could grant much more freedom of choice to each individual and withstand the resulting greater diversity among its members without degenerating into chaos. Thus, better information handling for administrative purposes could well result in a substantially better quality of education.

Characteristics of Time-Sharing Systems

The technique of time sharing (5) amounts, in essence, to the rapid time-division multiplexing of a computer among several users each operating on-line from a suitable terminal. This technique dates back to the early fifties and has been used in the Sage Air Defense System, in several airline reservation systems, and in various other special-purpose systems. In these so-called "dedicated systems", each user is limited to performing certain special actions in conjunction with a common data base. What is new in the various time-sharing systems developed since 1960 is that they allow users to do simultaneously totally unrelated work, thereby giving each of them the illusion of having a whole computer at his disposal. Still, the kind of work that a user is allowed to do varies considerably from system to system. Some of them make the entire core memory available to each user, although at different times, others partition the memory among the simultaneous users; some of them allow the users to program in only one particular language, others can be hosts to compilers for arbitrary languages. Thus the capabilities of existing time-sharing systems and the applications for which they are appropriate vary over a wide spectrum, and one should be wary of any comparison between them based on one or two characteristics. For instance, the number of simultaneous users that can be served by a system depends on who the users are and what they do, which in turn depends on what the system inherently allows them to do.

We shall discuss below various characteristics that are essential to achieving the goals outlined in the preceding section. None of the time-sharing systems currently in operation possesses all of the characteristics that will be listed as essential. Yet, all of these systems have proven to be useful for various purposes, at the very least as experimental tools for investigating the role of time sharing in specific areas. It should be stressed, on the other hand, that a system which is adequate for experimental purposes may be totally inadequate for every-day use on the part of a large community, and that the characteristics of a system influence greatly the types of applications that will be developed for it. People tend to shy away from applications for which the available tools are inadequate or just plain inconvenient.

The characteristic of a time-sharing system first observed by a user is its response time. Interestingly enough, this is the only characteristic that depends on the fact that the computer is "time shared" among several users. The response time may be defined as the time elapsed between the issuing of an instruction and the completion of its execution on the part of the computer; the time required to print out or display results is not included in the response time. Clearly the response time consists of two parts: the time that would be required to execute the instruction if no other users were present, and the delay caused by the fact that the computer must also execute instructions issued by other users. Thus the response time is a random variable which can be expected to vary over a broad range even for the same instruction; it is far from clear how one should represent it in a meaningful way or set specifications for it.

There seems to be two distinct ways in which a person can work comfortably in conjunction with a computer, (or for that matter with any other machine or even with another person), namely on-line or off-line. Working on-line implies that a person's attention is focused continuously and exclusively on one task, and therefore his mind can neither turn to anything else or rest while waiting for the execution of instructions on the part of the computer. Working off-line implies that a person's attention is divided in time among two or more tasks, so that the person can turn his attention to another task after issuing an instruction to the computer, and then look at the results of the execution later on at his convenience.

The length of time that most people can remain mentally and physically inactive without becoming impatient while waiting for the computer to execute instructions seems to be very short, about ten seconds. On the other hand, it is hardly worthwhile turning one's attention to anything else for less than a minute. Thus, response times in the tens of seconds are very frustrating. My impression, based on personal experience and on observation of other people, is that a person, when faced with response times in the tens of seconds, reorganizes his work to make the response time comfortable for either on-line or off-line operation. He may achieve this by breaking a computation into several smaller steps, by having the computer indicate periodically how far the computation has gone, by having the computer form several computations in sequence (even if some of them may turn out to be unnecessary), and sometimes by choosing a totally different approach to the problem under consideration. It is also my impression that it does not matter much what the actual response time is, provided it is smaller than the "threshold of impatience." Thus, the frequency with which this threshold is exceeded seems to be a more important system characteristic than the average response time.

Two conclusions can be drawn from these comments about response time and working habits of people. In the first place, the response time of a time-sharing system should be, most of the time, smaller than ten seconds for computations requiring up to a couple of seconds of computer time. In the second place, the two modes of operation, on-line with rapid man-machine interaction, and off-line with delayed inspection of results, should both be convenient and economical; they should also be compatible with each other so that a user may switch from one mode to the other at his convenience. In other words, time-sharing systems should offer optional batch-processing service under user control from a remote terminal, with the results being automatically recorded in the user's file for inspection at his convenience. Mixing off-line batch processing with on-line interactive service has the added advantage of permitting the utilization of computer time that might otherwise be wasted.

The value of a time-sharing system depends largely on the variety of facilities that it makes available to the users' community for work in diverse fields. The wealth of these facilities can be expected to grow with time as user groups develop special facilities for their own field. Crucial to this process, however, is the system's ability to host new facilities as a routine matter, that is, with a minimum amount of effort and without disruption of the system's normal operation. Systems differ greatly from one another in this respect. For instance, some of them can be used only for a single application, such as computer-aided instruction, and some of them allow programming in only one language, while others can be host to arbitrary sub-systems.

The generality of a time-sharing system, that is its ability to host arbitrary sub-systems, is important for at least three reasons. In the first place, it would be inconvenient and uneconomical to have different systems for different applications because of the resulting duplication of equipment and other facilities. In the second place, it is impossible to identify from the beginning all the facilities that are necessary for a particular application, and it may be impossible to add later on facilities that have been forgotten unless the system has a general-purpose structure. Finally, and most important, human activities are not compartmentalized, particularly, in a university; rather, they strongly interact and build on one another. Use of separate systems would set up artificial barriers which would in turn defeat the main purpose of facilitating intellectual communication throughout the community. This point was already emphasized in the preceding section in conjunction with computer-aided instruction.

Passing now to a different aspect of time-sharing systems, it is essential that the interface to the users be convenient from an intellectual standpoint, not just from a physical standpoint. Each user should be able to deal directly with entities and operations of interest to him, rather than with those that happen to be significant to the units comprising the computer installation. In other words, the characteristics of the hardware should be hidden from the user. For instance, the user should not be concerned about the size of the core memory, or with where programs and data may be stored at any given time. The technique of "program segmentation" (6) used in some of the systems (7,8) currently under development, should be of considerable help in this respect.

Segments are entities which preserve their individualities within the system at all times. Each user is free to decide which programs, sub-routines, or groups of data he wishes to treat as segments, and can assign to them names of his own choosing. Of course, there have to be limits to the number of segments that a user can execute within a single process, and to the size of each individual segment. However, they can easily be made so large (e.g. 2^{18} segments of 2^{18} words each), that no significant restriction is placed on any application. Each segment can make reference to a word in another segment by citing the segment's name and the position (symbolic or numeric of the word within the segment, regardless of where the segment may be physically located at the time. Segments are automatically retrieved and transferred to core memory when needed, and any necessary links to them are set up, also automatically, at execution time. In other words, all memory management is done automatically, and any segment can be used at any time without advance planning.

Secure, yet flexible control of access to a system and its contents is another very important requirement. In the first place, it is essential that no user be able to damage, either by accident or by willful act, the files of other users or in any way interfere with their work. Aside from any question of privacy, it would be practically impossible to debug a program if it were possible for it to interact by accident with other programs. Positive identification of each user through some sort of personal password is essential, but this is only one of many precautions that must be taken by the system. Maintaining security in a time-sharing system is a complex problem from both the operational and the design standpoint, which is beyond the scope of this discussion. It is worth pointing out, however, that the programmers entrusted with maintenance of the system are bound to make occasional mistakes which may invalidate the system's security. Therefore the system must possess several lines of defense against intruders. It is also important that the system include facilities for automatically tracing security violations. Experience has shown that vandalism within a time-sharing system and the forging of user accounts are to be expected in universities as well as elsewhere.

While it is essential to prevent unauthorized use of files, it is equally essential to facilitate as much as possible intellectual communication between people. Specifically, a person should be able to authorize other people to use one of his programs or groups of data, and yet prevent them from modifying it, or copying it. Also, it should be possible to specify the purpose or the circumstances for which the authorization is given. For instance, some data may be made available to a person for statistical purposes only, or without allowing him to make correlations with other data. This type of control is needed in conjunction with programs for computer-aided instruction to which students and instructors should have different types of access.

We have so far spoken about control of access only in conjunction with users' files. Similar considerations and needs apply also to the supervisory programs of the system. The point is that it should be possible to adapt the system and its appearance to the needs of different people. For instance, the system should appear very simple and easy to use to a beginning student, while the system's sophistication and its resulting complexity should be apparent to the expert user. For this purpose different users may be given access to different supervisory programs or to different entry points of the same program. For instance, different users' requests may be scheduled by different algorithms. It is worth mentioning in this respect that the implementation of such features requires, in practice, that the pertinent supervisory programs consist of independent entities in the sense of the segments mentioned above. In other words, the technique of program segmentation seems to be essential to adapting the system to individual user needs.

The value of a time-sharing system depends largely on its ability to act as a repository of the knowledge of the community. Thus, the users must be willing to entrust the system with their work, and the system must be able to protect user's files against mishaps that may damage them. For this purpose, the system must automatically copy all users' files on magnetic tape shortly after they are generated or modified, and procedures must be established for reconstructing the contents of the system's mass memory from the magnetic tapes in case of mishap. The complex of facilities and procedures for protecting users' files, often referred to as "file backup", is essential for any time-sharing system serving remote users.

A time-sharing system presents some new problems of administration which place additional requirements on the system itself. Of the various issues involved two affect directly the individual users, namely the decentralization of resource allocation and the metering of equipment use for charging purposes. The administrative actions involved in the operation of a time-sharing system have to do mainly with the changing of individual passwords and the allocation and reallocation of resources whether dollars, computer time, or space in the mass memory of the system. Experience has shown that long frustrating delays are unavoidable if these actions can be taken only by personnel of the central administration. Thus, it is essential that the users themselves or their immediate supervisors be able to perform these administrative actions without the intervention of the central administration. On the other hand, the tables where such data are stored are extremely sensitive, and unauthorized access to them must be prevented. Thus, we encounter here too the need for secure access control. Another important issue is the basis on which users are charged for the use of the system. At present, charges are based almost exclusively on the central-processor time used, with the result that some users have to pay for a lot of peripheral equipment they do not need, because such equipment has to remain idle while the central processor is being used. Since a time-sharing system serves many people at the same time, the various units in the computer installation can be utilized more uniformly over time. Thus, it may become economically possible to charge each user only for the system resources that he actually consumes. This requires, however, that the system be able to measure continuously the use made by individual users of each unit of the computer installation. Of course, the algorithm for computing the overall charge should be a peripheral part of the system, and easily modifiable by the system administration. On the other hand, the facilities for metering the use of equipment must be planned from the start in the design of the system, to insure that the process of metering will not become an intolerably large system load.

Finally, a few words should be said about terminals, although their characteristics are not intimately interwoven with those of the system. Current time-sharing systems use largely teletypewriters as terminals. The reason is very simple. Teletypewriters are mass-produced for other purposes, and therefore they are available in quantity at a relatively low cost. On the other hand, typewriters are clearly inadequate for man-machine interaction. In the first place, the rate at which they can display text is far too low. In the second place, the fact that a teletypewriter can not display graphical information is a very serious limitation. Of course, display terminals are currently available which are more than adequate, but they are far too expensive for general use. A display terminal recently developed at M.I.T. (9) appears to meet both technical and cost requirements for general use. It includes a storage display tube capable of displaying 4,000 characters and circuitry for generating them at the rate of 120 characters per second when transmitted through a 1200 bits/sec. telephone line. It can also display line drawings of arbitrary complexity. It can be constructed at present for a little less than $10,000, and a much lower cost could be achieved in mass production. Industrial interest in display terminals is growing, so that adequate devices should be available at reasonable cost in the near future.

Cost Consideration

No discussion of the role of time-sharing systems would be complete without some consideration of their cost. We shall use as a specific example the Compatible Time-Sharing System (CTSS) which has been in daily operation at M.I.T.'s Project MAC for the last four and one-half years. This system employs an IBM 7094 computer installation, a fact which makes it easy to compare its cost with that of a typical batch-processing system.

A question very often asked about a time-sharing system is: How efficiently is the computer utilized? In the Compatible Time-Sharing System the central processor is executing user instructions 40-50 % of the time. Thus, the utilization of the central processor is approximately the same as in a conventional batch-processing installation. The central processor is executing supervisory instructions (including those pertinent to switching its attention from one user to another) 10-15% of the time. The rest of the time (40-50%) the central processor is idle waiting for data transfers to and from the core memory. In order to utilize this idle time one would have to have more than one program (ready for execution) in core memory at all times. In such a case, however, the memory occupied by programs that are not being executed would be in effect wasted, and core memory is almost as expensive as the central processor. Furthermore, the cost of the central processor is less than 20% of the cost of the whole installation used by the Compatible Time-Sharing System, and about 35% of the cost of a normal batch-processing installation. Thus, it is misleading to equate utilization of the central processor with efficient operation of the whole computer installation.

The fact that the Compatible Time-Sharing System and a conventional batch-processing system utilize the central processor to about the same extent does not mean that the two systems can provide computer service at the same cost. The computer installation needed for the time-sharing system is much more expensive than that normally employed for batch-processing, as indicated in Table I. The equipment in the left column of this table is that pertaining to a normal batch-processing installation; the additional equipment needed for time-sharing is listed in the right column. The costs are given as fractions of the total cost of the batch-processing installation. Thus, the total cost of the installation needed for time sharing is 86% greater than the cost of the batch-processing installation. The first additional item is a second core memory for storing the supervisory program and for providing necessary buffer space. The second item includes all the equipment needed for multiplexing the system terminals (including the display terminals). The third item, the largest in cost, includes the various mass memories in the system and the related channels.

Users are charged for central processor time at a rate of $300 per hour for the day shift, Monday through Friday, and $200 per hour at all other times. They are also charged for storage in the disk file at the rate of 20 cents per month per record (432 usable words), and for drum storage of personal file directories at the rate of $30 per month per record. An access or registration fee of $20 per month is also charged to each user to defray overhead charges and to discourage inactive users from remaining registered. The total income from these charges is approximately equal to the cost of operating the time-sharing system, which totals approximately $90,000 per month. This operating cost includes the rental of the computer installation at approximately 50 per cent educational discount, personnel salaries, power, supplies, and various other expenses pertaining to the operation of the system. Charges for terminals are billed separately and vary depending on the type of terminal and its location. The on-campus charge (including data set and communication costs) for an IBM 2741 teletypewriter is $131.50 per month and for a model 37 Teletype is $122.50 per month.

The maximum number of simultaneous users is usually set at 30 (it actually varies automatically depending on total system load). The ratio of total terminal time and total computer time turns out to be also in the neighborhood of 30. Thus, an hour of terminal usage during the day shift costs approximately $10. An additional $1 per hour would cover the terminal cost for student use.

On this basis, one terminal hour per week of use on the part of a student would cost about 16% of his tuitions (assumed to be $1,000 per 15-week term).

Spending 16% of each student's tuition on computer usage is not in itself unreasonable. On the other hand, significant use of computers in education would require at least 10 terminal hours per week per student rather than 1 hour. Thus, a reduction of cost by a factor of 10 is needed.

A cost reduction by a factor of 10 is not very large when one considers the progress that has been made in the computer field during the last 10 years. The higher speed of the newer computers together with better system design should yield a cost reduction of a factor of 3 within a year or two. A reduction by a factor of 10 may well be achieved within 5 years,and almost certainly it will be achieved within 10 years. On the other hand, it will take at least 5 years to rearrange our curricula to take advantage of computers in a substantial way, and to develop the necessary software. Thus, educational developments rather than cost may end up to be the limiting factor in the exploitation of computers in engineering education.

Conclusions

The broad goal of a time-sharing system, in engineering education as well as elsewhere, is to help individuals in their intellectual work, by acting in their regard as a knowledgable and skillful assistant. An essential aspect of this role, particularly with respect to education, hinges on the system's ability to facilitate intellectual communication and interaction between people, and to act as a repository of the knowledge of the community. Of particular importance in this respect are the convenience and flexibility with which access to private files can be controlled, and the reliability with which the files are protected against accidental or willful damage.

Time-sharing systems can be very powerful tools in engineering education, and can free it from many of the present constraints which are sources of frustration to teachers and students alike. The potential educational benefits are great indeed, particularly, with respect to adapting instruction to the individual needs and desires of students. On the other hand, the exploitation of time-sharing systemshas barely begun, and many problems, technical and educational, must be solved before significant benefits will be realized. In the first place, none of the current time-sharing systems is technically adequate in all respects. In the second place, the cost per student-hour must be reduced by at least a factor of 10 before widespread use on the part of studentscan become economically feasible. Moreover, a "bare" time-sharing system as it may be provided by a computer manufacturer, will be of little use except for programming purposes, until instructional and other special subsystems have been developed. The task of developing such subsystems is a major one, at least as difficult and time-consuming as the preparation of atotally new set of textbooks for all subjects of instruction. Finally, extensive use of computers in engineering education will undoubtedly lead to major changes in what we teach, in the way we teach it, and in the structure of our curricula. Thus, we will have also to face the organizational and human problems inherent to such major departures from past practices, and find adequate solutions to them.

The technical and financial problems that we are facing today with respect to time-sharing systems look large and perhaps overwhelming. Yet, technical progress and cost reductions have been so rapid in the past, that we can expect these problems to become of secondary importance within a decade. Thus, the extent of our progress in solving the educational problems will probably be 10 years from now the limiting factor in the exploitation of computers in engineering education. The fact that time-sharing systems are at present too expensive and not fully adequate should not discourage us from tackling immediately the many educational problems which undoubtedly will turn out to be more difficult as well as crucial.

TABLE I

Basic Equipment	Fractional Cost	Additional Equipment	Fractional Cost
Central Processor	.35	Core Memory	.23
Core Memory & Multiplexor	.30	Transmission Control & Two Channels	.23
Two Channels & Peripherals	.35	Two Channels, Three Drums, & Disk File	.40
	1.00		.86

Distribution of equipment cost in IBM 7094 installation employed by the Compatible Time-Sharing System.

The place of time sharing

[1] H. Sadman, W. J. Erikson, and E. E. Grant, "Exploratory Experimental Studies Comparing On-line and Off-line Programming Performance," Communications of the ACM, Vol. 11, p. 3, January 1968.

[2] J. Weizenbaum, "Contextual Understanding by Computers," Communications of the ACM, Vol. 10, p. 474, August 1967 - See references for other pertinent work.

[3] J. Moses, "Symbolic Integration," Technical Report MAC-TR-47 (Thesis), Project MAC, M.I.T., December 1967.

[4] J. F. Cogswell and D. P. Estavan, "Explorations in Computer-Assisted Counseling", TM-2582/000/00, System Development Corporation, Aug. 6, 1965.

[5] R. M. Fano and F. J. Corbató, "Time Sharing on Computers," Scientific American, Vol. 23, No. 3, p. 128, September 1966.

[6] J. B. Dennis, "Segmentation and the Design of Multiprogrammed Computer Systems," Journal of the A.C.M., Vol. 12, p. 589, October 1965.

[7] F. J. Corbató, V. A. Vyssotsky, "Introduction and Overview of the Multics System," A.F.I.P.S., Conference Proceedings, Vol. 27, Part 1, p. 185, 1965.

[8] W. T. Comfort, "A Computer System Design for User Service," A.F.I.P.S. Conference Proceedings, Vol. 27, Part 1, p. 619, 1965.

[9] R. H. Stotz and T. B. Cheek, "A Low-Cost Graphic Display for a Computer Time-Sharing Console," Proceedings of the Eighth National Symposium on Information Display, p. 91, May 24-26, 1967.

[10] P. A. Crisman, Editor, The Compatible Time-Sharing System, Second Edition, The M.I.T. Press, Cambridge, Mass., 1965.