This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

UCRL - 74915

PREPRINT

CONF-730964--1

LAWRENCE LIVERMORE LABORATORY

University of California/Livermore, California

NEW COMPUTER TECHNOLOGY TO BE APPLIED
TO DECISION MAKING PROCESS

Sidney Fernbach

August 1, 1973

NOTICE

This report was prepared as an account of work sponsored by the United States Government. Neither the United States nor the United States Atomic Energy Commission, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness or usefulness of any information, apparatus, product or process disclosed, or represents that its use would not infringe privately owned rights.

Conference of the Society
For Management Information Systems

RECEIVED BY TMO DEC 12 1973

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

NEW COMPUTER TECHNOLOGY TO BE APPLIED
TO DECISION MAKING PROCESS

Sidney Fernbach
Univ. of Calif., Lawrence Livermore Laboratory

I. STATE OF THE COMPUTER INDUSTRY: "FOURTH GENERATION"
HARDWARE

A. Central Processors

We have seen three computer generations go by in the last twenty years. These were differentiated primarily by the components utilized as basic ingredients of the hardware system, going from vacuum tubes to transistors to integrated circuits. Now we are in the highly integrated circuit era, dealing variously with LSI (Large Scale Integration), MSI (Medium Scale Integration), and SSI (Small Scale Integration). These are distinguished by putting 1000 circuits on a chip for LSI, 500 for MSI, and 100 for SSI.

It is not clear that we will be able to talk about future generations in terms of components again. The IBM 370 series is sometimes called the 3-1/2 generation. Perhaps the 380, if there is such a thing, would then become 4th generation. Furthermore, dependent as we are on IBM's moves, if there would then follow a 390 series, we would call that the 5th generation. In actual fact, we are beyond that now. Some of the smaller computers being built by independent companies use more advanced components than does IBM or some of the other major manufacturers. Indeed, Control Data Corporation has been using 2nd generation components (discrete transistors) for its top of the

line computer, the 7600. This system competes with the highest end of IBM's 3-1/2 generation. As a matter of fact, it is our understanding that CDC has on its drawing boards another even more advanced system using 2nd generation components.

Because of these discrepancies in defining generations of computers, those of us involved in applying the high end of the various lines to our problems now resort to categorizing computers by class. We call the high end of the IBM 360/370 as well as the CDC 7600 a Class IV computer; those now being produced to operate at higher performance, Class V computers (such as TI's ASC(4X), CDC's STAR and Burroughs ILLIAC IV); and finally those now on the drawing boards, but intended to provide higher performance still, as Class VI systems. Much of the philosophy, architectural design, and family similarities are passed on to computer systems of lesser performance produced by the same manufacturer. As yet, it does not seem that the many architectural innovations that appear in the minicomputers have found their way up to the giant machines. Perhaps they never will, except in portions of the large systems.

During the 3rd generation, the family concept was developed to its highest point. There has been some deterioration in time, mainly because the manufacturer now takes advantage of newer ideas or developments as they occur. It is virtually impossible as well as undesirable to bring out a new family every year; it is more feasible to spawn additions. Within the family a

wide range of performance capabilities has proved to be possible. As advantage is taken of the newer technology, performance improves and costs drop. This is due primarily to providing more circuits on a chip and shorter distances for signals to travel. The change from IBM 360 to IBM 370 indicated very clearly a performance improvement of approximately a factor of three with no change in price. One keeps looking for greater cost effectiveness as the higher technology components are introduced. See Figure 1, "Cost Effectiveness of Computer Families."

If we believe what we hear, the new Armdahl computer is reportedly more capable than its 370 equivalent at roughly the same price, again by a factor of the order of three. Perhaps we are reaching some limitations in performance such that the next improvement in cost effectiveness will show up as a reduction in price for the same performance. This, of course, is also true now if one compares, say a 370/155 with a 360/65. The availability of funds and the psychology of the user (and seller as well) is such, however, that one buys a 370/165 instead of the 370/155 when one is moving from the 360 to the 370 system. I am not singling IBM out for these remarks as providing the only example of using the counting system for a psychological advantage. Their system merely represents a better known set of numbers. All manufacturers take advantage of these selling gimmicks.

B. Memory Systems

Since the ferrite cores became available in the mid-fifties, they have become the most common ingredient of computer systems and remained so through the various generations until the present. Size of cores and hence speed have changed over the years so that cores approximately 10 mils in diameter with submicrosecond flipping speeds are in very common use today. Furthermore, the prices have come down from dollars per bit to cents per bit. Where a few thousand words of memory was all one could afford in a system in the 50's, we now use hundreds of thousands of words. As indicated above, general performance of computer systems at lower costs has improved many fold over the past twenty years. In no small way, increased memory size and improved core speed have been instrumental in improving throughput, which in some respects is more important than the raw speed improvements cited above.

We are now in what appears to be another revolutionary period in memory development. Not since the magnetic core was first introduced has the great switch to new components been so rapid. The new technology involves the semi conductor memory, i. e., memory on a chip. IBM again has taken the lead in producing these new memory systems for its 370 line. Potential performance is greater than that of cores and packing densities are inherently greater, primarily because more and more bits can be put on a chip. Prices, although comparable

with cores today seem as if they will tumble faster than that for the cores. Four thousand bits on a chip measuring 150 mils on a side is becoming fairly routine now. With such densities, a few pennies per bit allows for head-on competition with cores. At these same prices, the more recently developed N-channel MOS (Metal Oxide Semiconductor) memories are beginning to appear with speeds in excess of the rival cores.

The more expensive bi-polar semiconductor memories are exhibiting fantastic speeds, in the 10 nanosecond range, allowing for dramatic speed-up in performance (if one could afford to populate a computer memory with such devices). In actual fact, of course, the cache memories commonly used in the 3-1/2 generation systems are bi-polar, but with somewhat less performance than that mentioned above.

C. Peripherals

Peripheral devices have changed more in type and variety over the years than other portions of computer systems, but not as dramatically. Speeds of processors and memories have increased by factors of thousands whereas speeds of tapes and other magnetic peripherals have increased only by factors of hundreds. Card equipment and printers have improved by factors of tens. Interestingly enough, despite such minimal improvements, these devices are still with us.

The most important peripheral devices are those intended for main memory backup. Tape was the most significant medium during the early years; but because of the impossibility of providing random access and because of poorer reliability, the disk or drum has been gaining a more prominent position in the system structure. The diskpack system especially is relegating the magnetic tape from its secondary role to a tertiary one. Currently, the highest state of the art is exhibited by the IBM 3330 and its equivalent, manufactured by those who produce equipment "plug compatible" with IBM systems. A single spindle consisting of a stack of 10 platters will store 800,000,000 bits of information. Furthermore, these bits are accessible within tens of milliseconds at many megabits per second rates. There will be continual improvement in this type of device.

Magnetic drums have always represented a problem for computer users. The cost per bit of drum storage is intermediate between that of main memory and disk. For the sake of argument, one can estimate the cost of main memory as one cent per bit, fixed head drum memory at one-tenth cent per bit, disks at one-hundredth cent per bit. Early drums had moving arms and heads. Today's drums generally have a fixed head on each track. This is so because it is important to provide as short an access to information as possible. In a fixed

head device the only time involved is the rotational time of the surface. Unfortunately, the cost of heads is so high that the drum must be kept to a minimal size, hence providing a relatively low capacity in order to maintain the price at the level mentioned above. Thus, the drum offers extensive memory, but not sufficient to make it as desirable as a disk. On the other hand it offers better access time. It seems to be an attractive nuisance, since its inclusion in a system also requires special software for yet another level of storage. The drum is clearly on its way out, but it is not certain yet what its replacement could or should be. This will be considered later.

Other forms of supplemental memory have been proposed and even built, but none have been widely accepted. There have been wide magnetic strips, such as appeared in IBM's Data Cell and RCA's RACE; video tape systems such as the AMPEX terrabit memory and the IVC recorder; photographic recording on IBM's photo-digital store (Model 1360), Precision Instrument's laser recording on long wide metal coated mylar strips, etc. Even though some of these were, or are, capable of storing up to trillions of bits, they did not catch on--they never became part and parcel of complete systems. Some are still available today, and as a matter of fact, there is a revitalization of very large archival memory developments.

Magnetic tape systems have improved over the years,

perhaps not quite as they should have. Generally the tape reel that was used twenty years ago is still being used today. There may be more information stored per inch, the tape may be driven faster, the reel may be self-loading, but the entire device is bulky, noisy, and as unreliable as ever. Tape format changes have been coming relatively fast in the past few years, so there is nothing but a moving de facto standard. So long as tape can remain a tertiary storage medium this may not present much of a problem, but as a medium for exchange of information this leads to many problems.

Printers have improved somewhat, but have remained slow and expensive. Because they have not kept pace with improving computer performance we now have multiprinter installations. The impact printer has represented a fine device in the past because of the fact that the paper used was inexpensive and the quality of print could be brought to fairly high standards (at moderate speeds). Non-impact printers have recently been introduced which allow for greater speed, lower capital costs, high resolution (to the extent that good graphic printing can be produced), but with higher paper costs.

There is an increasing use of microform output also. The so-called COM Units (Computer Output Microform) are very useful devices if one wishes to conserve storage space. One can put computer output on 16mm, 35mm or 105mm film (the latter

usually being cut up into short strips called microfiche), and store these in canisters or small file cabinets. Enlarging, viewing, and duplicating equipment is available in large measure to handle this medium. Generally, COM output is fast to produce and inexpensive, but it is not easily viewed. Also it is more difficult to make easily visible notes as one does on paper.

One of the major peripheral devices that has proliferated over the past 10 years is the computer terminal. Ever since time sharing or interactive use of the computer became prominent, the rush has been on to develop a better, less expensive keyboard or display device to allow individuals to communicate with the central computer. These are varied but fall into several classes, such as keyboard/hard copy, keyboard/soft copy, alphanumeric versus graphic, with or without light pen or cursor, and with intelligence or not.

D. Software

Operating systems have been developing over the years, becoming more and more complex with each year. It is getting to be just as doubtful that all the bugs in these systems will be found as it is that all hardware bugs are found during the lifetime of a computer system. It is difficult to see how this can go on without some structure for developing software systems being

implemented. There is talk of teaching "software engineering". It is hoped that someone discovers what this is so that it can be taught.

Most operating systems are fairly modular, so that more functions can be added as needed. The chief attributes of these systems are that they will take care of the basic needs of the user primarily in a batch mode. In some instances interactive use of the computer is permitted. One of the missing ingredients is security. Most existing systems which are openly available to remote terminals can be tampered with so that the privacy of one's files may be invaded. Also lacking in most systems is the possibility of making performance measurements, auditing procedures, integrity study measures, and the like.

After a great proliferation of languages, there seems to have resulted a rather quiescent period. There are new languages being developed, but they are not being pushed against the old standbys. Also, some of the newer languages are being designed for specialized purposes. COBOL and FORTRAN are predominately used in American computer systems. ALGOL still is widely used in some places in this country, but more so in European countries. It is not likely that this situation will change much despite the fact that some features of the newer languages are better.

E. Networks and Communications

One of the most significant computer related developments in recent years is the concept of networks. About six years ago someone in the Johnson administration dreamed up the phrase "networks for knowledge". It had such a good ring to it that the president's office immediately asked for comments from all agencies of the government concerning possible implementation of such a thing. Nothing much came of it because no one knew quite what it meant. There is more significance to networks for knowledge today, in that working networks have been put into operation where many of these are in the university environment.

A computer network may be considered broadly as a group of computers tied together and shared by numerous users. The most widely known network is that developed under the influence and support of ARPA (Advanced Research Projects Agency of the Department of Defense). Most of the nodal points are computer centers at large universities scattered throughout the United States. There are some nodes located at government-supported laboratories as well. Each facility has a small, intelligent computer to interface its own system to the network. Communications are accomplished through the use of high speed (50 kilobit/sec) lines. This network has been in operation for several years now and has shown the way to improved network

design, for developing protocol in communications, and, in general, has shown the feasibility of linking a variety of systems. There have been numerous discussions about forming additional networks of this nature, mainly for both institutes of higher learning and government use.

Smaller networks linking numerous terminals to several computers exist in greater abundance, but these are rather simple in comparison to the ARPA net and not difficult to implement. They do get extensive and effective utilization, but in a sense are no different from a single computer radiating out to a large number of users at terminals remote from the central computer (intelligent or otherwise).

The most impressive networks that are in use are the ones being used for "real" work, such as the airline reservation systems, point-of-sales systems, and industrial networks that tie together different plants of a given manufacturer, even world wide. There is no doubt that networks will expand and be more widespread with time. It is not clear yet how much information one must or should pass over the communication lines. Communications are expensive and do not seem to be headed for a decrease in cost. The sheer volume of extensive communications of data from centralized data bases could be staggering and result in chaos. Much must be done to determine the best strategy of minimizing the quantity of data to be moved in a network.

II. WHERE ARE WE GOING: CLASS V OR VI COMPUTER SYSTEMS

A. Computer Systems

Considering those computers now in existence as Class IV or "4th generation" if you will, what are we going to see in the near future? As mentioned earlier, three major systems designated as Class V are in various stages of completion. Two of these systems are referred to variously as "streaming", "pipeline", or "vector" machines, the third as a "parallel" processor.

Generally, most computers of the past have been sequential (or scalar) processors. This means that instructions are executed sequentially, each result being operated on or stored for future needs in serial order. Because of the finite size of components, mounting boards, and wire lengths as well as the time delays in sending signals through these devices, the ultimate speed of sequential computers is limited. However, computer speeds are still improving without major architectural changes, since the dimensions of the most recently developed components are decreasing and greater densities are being made possible through technological advances in the assembly of component structure.

On the other hand, factors of improvement were measured as decimal orders of magnitude in the 60's and will be more like orders of magnitude in the 70's. Hence, other changes

are being made in the architecture of computer systems to allow for operations to be performed in parallel, thus enhancing overall performance. The techniques currently being implemented are of three varieties. The first is a multiprocessor system which allows several processors to share a common memory and independently process data pertaining to a common problem or in a multiprogrammed manner. The second is a parallel processor consisting of a large number of small processors which may have independent memories, share a common memory, or both. These small processors operate in a lock-step mode; that is, each and every one carries out the same operation on its own operands at the same instant of time. The third technique is the "pipeline" processor which consists of a segmented arithmetic unit which allows sequential stages of instruction execution to be carried out in parallel on different sets of operands as they flow through the "pipe" pairwise.

Some steps have already been taken in Class IV systems and even Class III to permit multiprocessing. Univac has advocated the use of triple processor systems since the days of the 1107. IBM has had dual processors for a considerable length of time as standard equipment. Some special purpose equipment, such as that built for FAA, for example, used more than two processors in a system. CDC also has paired processors in several systems in the 6000 and 7000 series of machines.

These are but examples; others have done likewise. In most instances, however, the multiprocessor is programmed to do designated tasks in a multiprogramming sense, rather than participating in the solution of single, large-scale problems.

It is not at all clear how many processors can be made to share a single work load and effectively gain performance. This may be a rather small number as low as four. Such or similar systems are now being planned as potential Class VI machines. In the scientific environment, there is no doubt that attempts will be made to multiprocess single large jobs when these machines appear. Very little effort currently is being expended to study this type of problem organization, programming, or its implementation. In the case of Management Information Systems, it is highly unlikely that these systems would be programmed to single large jobs; rather, they would be used as multiprogrammed devices to achieve n times the performance of a single machine for n distinct jobs.

The parallel processor was conceived several generations ago, but is only now making its appearance in the form of ILLIAC IV. This machine was conceived by Professor Slotnick of the University of Illinois, built by the Burroughs Corporation and funded by ARPA. It is now located at the Ames Laboratory operated by NASA in Mountain View, California. The computer is virtually complete, and problems are now being debugged on

the machine. Since this powerful tool is being made one node of the ARPA net, additional peripheral devices are being added to the system. This may delay the widespread utilization of the system. Since the machine has not been successfully tried and proven, it is difficult to assess the performance of this computer architecture.

The NASA system consists of 64 parallel processors each with several thousand words (16K bytes) of memory. These processors together are capable of cooperative action on individual problems. The system was intended to solve mathematical problems of such a repetitive nature that 64 identical calculations could be performed simultaneously in a consistent fashion. If only a small fraction of these are used, the effectiveness of this organization drops substantially. Burroughs is still investigating this structure as well as modifications of it for future scientific markets. Additional sales may well depend upon the success of the system installed at Ames.

So far, as business data processing is concerned, the parallel organization could be quite effective. Since little or no exploratory work is going on in this area, it is highly unlikely that this type of machine will be used in this area for some time to come.

The "pipeline" structure is somewhat similar to ILLIAC IV in its parallel processing capability. Both Control Data Corpora-

tion with its STAR and Texas Instruments with its ASC (Advanced Scientific Computer) have built systems based on this architecture. In either case, the most effective utilization involves setting up the data so that it can be streamed through the arithmetic unit for a particular operation, the results being put away, and then the next "vector" stream being set up to be operated upon. If the stream is not very long effectiveness is sacrificed. If the stream is long, speeds equivalent to 50 million instructions per second can be obtained. Both CDC and TI have sold such systems and are building enough of them that valuable experience with this type of architecture for both scientific and commercial use should be available within the next few years. At the time of this writing, only TI has delivered a system of sufficient size to test thoroughly.

It is difficult to guess whether or not the parallel concept will catch on in the commercial world. Where ultra high performance is required, the parallel concept may be the only way to beat the restrictions imposed by the finite speed of light.

Of course, the march of technology is allowing for higher-speed sequential machines to be designed also. The late 70's will see their appearance. It is more likely that these will become the workhorses of the 80's for scientific and business data processing. They will be designed such that the heavy investment in current software can be utilized, permitting for a

more gradual transition to additional software goodies such as security and audit features. These systems may become the heart of the computer utilities that we have heard so much about in the past. That is, they may become the heart of the network and the tremendous data banks.

The main memories for these systems will, no doubt, be of the semiconductor variety. They will be substantially large in capacity, with eventual costs in the under one cent per bit range. In addition, speed will be somewhat greater than now on the average. Those systems which are designed to be ultra fast will have the more expensive bi-polar memories, perhaps in more limited quantity. As in the past, there may be hierarchies of main memory to allow for large capacity at lower overall cost. With the IBM-announced endorsement of paging, it is logical to assume that paging will be a standard for future systems. It allows the user greater flexibility in addressing data, at some slight expense in response time.

B. Peripherals

The magnetic disk pack system seems destined to be the main random access storage file for some time to come. Clearly, one can add platters to the stack, increase densities, rotate them faster, read and write tracks in parallel, and add more intelligence in the controller to improve on the three important features

of such storage devices: capacity, access time, and transfer rate. Each of these areas of improvement is currently being worked on. Prices for these improved species will not be appreciably greater than the prices for the current products. Because of these possible advancements in disk pack systems, it becomes obvious that fixed head drum systems may not be cost effective in future systems.

Attempts in the Laboratory to substitute optical devices for magnetic devices have not been very successful. Although it seems plausible that a higher density of storage might be achieved optically, improvements in magnetics have matched the performance in optics. The one bright area for the future seems to be holography. By using a laser beam to store information on special material, it may be possible to achieve densities impossible to reach with magnetics. The material problems are such, however, that it may take 5-10 years to get to the point of applying these techniques to computer storage devices. Once these holographic techniques are developed, not only will the data storage problem be significantly affected, but the entire computer structure could be drastically altered. This type of speculation will have to wait until the development seems more hopeful.

Other changes in peripherals seem to be coming gradually, and although they may change the way of dealing with input and output,

there is little to foresee that would be revolutionary. We expect more intelligent terminals, the increased usage of diskette or its equivalent (e.g. cassette), higher performance printers and microfilm units, greater emphasis on pictorial representation (with both color and sound), and a move towards networks.

C. The Minicomputer

While this progression to bigger and faster systems has been taking place in a more or less predicted fashion, a quiet revolution has been going on. This is the advance of the minicomputer. In sheer numbers being developed, the curve is rising at a staggering rate. See figures 2 and 3, "Population of Installed Computers." The economic impact has not been so great because of their low price. The cost of a small system today is no more than one-hundredth the cost of the first electronic computers with which the mini's compare in capability. Thus far, most of them have been put to work in data acquisition and control systems in industry and the laboratory. They are also being used in terminals, for local, small-scale data processing, and in schools for simple, low-cost time-sharing systems.

The highest level technology known today is going into these systems, making them both more capable and less costly than before. The number of manufacturers producing these

computers is also large, hence there is a variety of models each with its own software. In some respects, it appears to be quite a mess. On the other hand, each software system is relatively simple and is produced by a handful of programmers rather than an uncontrollable horde. It is far easier to put a mini into operation than a large computer. One might ask, can we not explore parallel or multiprocessor systems using mini's thereby getting more capability for fewer dollars and less effort to test these schemes. The answer is yes. Such experiments are being carried out.

The greatest advance in the minicomputer is now taking place. During the next decade it may result in a greater impact on data processing than the large computers discussed above. This advance is the utilization of the concept of microprogramming - first proposed 20 years ago and well before the time it could be most advantageously implemented. In the scientific computer world we went from special-purpose analog devices to general-purpose digital devices. Now we can go back to special-purpose devices but of the digital form. Business data processing grew to where it is now with the general purpose digital computer. Now it also can retreat to the special-purpose device. I should like to propose that, in the future, more cost effective data processing will be performed on small, special-purpose computers combined in a network sharing the requisite

data resources. The software will have been replaced by firmware dedicated to specific functions in each mini in the network. These computers will be simpler to program, more reliable, inexpensive enough to allow for adequate redundancy, and simpler to protect than large machines. Furthermore, it should be simple enough and inexpensive enough to provide for replication of systems such that national networks or utilities should not be necessary or desirable.

As yet, it is not clear how one would treat data bases. Should it become convenient to mass produce recorded diskettes, or readily and cheaply duplicate information stored on such media as tape in a cartridge or cassette, there may be no need to provide extensive, centralized data banks. I believe these developments will occur.

It is not being suggested here that large computers will disappear. There will always be larger problems than we can handle on whatever equipment we put together. Furthermore, the design of the programs for the special purpose computers and exploratory phase of new developments will have to be performed in larger, more general systems. There will continue to be real-time requirements that demand the highest performance attainable. A good example is that of weather predictions.

III. EFFECTS OF CHANGING SCENE ON MIS

A. Trends

Unfortunately, I am no expert on MIS and cannot speak directly to this point. As I have watched MIS grow, it seems to have done so in a modular fashion. Actually this is the only way in which it was possible to progress. A complete system, if there is such a thing, becomes a tremendous undertaking, much larger than any scientific problem we are presently solving. The software could become so involved that no checkout would be possible. On a modular basis success in creating a working system is possible.

The nationwide investment in software to date must be very impressive. No doubt there would be great reluctance to rewrite much of it, or even to make major modifications. Hence, the manufacturer who finds his greatest market potential in the commercial community also is reluctant to make changes which have any great impact on the user. Hence, the family concept and compatibility from generation to generation will continue to be stressed.

The software system will have to change somewhat to allow for the goodies that are now lacking: the guarantees of privacy, integrity, and security; the ability to provide audit tracks; the capability to measure performance levels; the capability of providing pictorial data, voice data, and so on. All aspects of the software systems will become more complex as more and more

capability is demanded by the user. Finally, the system will collapse under its tremendous weight and overhead.

Not to be overlooked is the possibility of having computer needs provided by service organizations through a network system. With these very large computer systems, notably the Class IV, V, and VI systems, it is conceivable that canned packages may be available to provide for the needs of the smaller commercial users. The larger users may still find it less expensive to house and operate their own equipment.

As mentioned above, there may be an alternative to the concept of a single processor with an ever-growing software system. This is the network of specialized minicomputers, each designed to carry out a specific task and communicate with a common data base. Such a system could be used very effectively in MIS as I view it. Every specialized function, of which there are many, could have its own dedicated "firmware" oriented system. For example, payroll, procurement, inventory, and so on, each would be designated as specialized functions. Rather unique, simplified programs could be implemented in read-only memory to carry out their functions. The data base could be shared with other functional units if desired or kept with the dedicated system. Incidentally, the management of the data base is another function requiring a special unit, as is the communications control, input and output

control, and as many others that may be needed to carry out the entire management operation.

B. Summary

Much of what has been presented here is indicative of our national trends in computing. To review and bring in the international scope at the same time, let me enumerate the following point:

1. There are but few new ventures in the competition for the high performance end of the computer market. Many of these look to IBM as the leader and hope to be compatible either software-wise or at least peripheral-wise. The Japanese are strong contenders in this market. A combination of European nations is being formed to compete as well. In the Eastern Bloc countries, there is talk of large systems but no evidence as yet of their existence.

2. The competition for the middle range of computers is very great internationally. IBM dominates this market, hence sets the pace for software and peripheral development.

3. The number and varieties of minicomputers is increasing at a phenomenal rate throughout the world. There is no question but that the number of manufacturers and number of systems built will continue to accelerate for some time to come. The microprogramming concept will be pushed to the

extreme and will find a place in mini structures that will allow these small systems to be used in all areas where specialized functions must be carried out.

4. In all these systems, because of rapid price decreases, memory size will increase. Predominently, main memory will be of the semiconductor variety.

5. Software is becoming more complex and will continue to do so. Most systems will be based on existing systems. Attempts to impose standards will continue to be made, some even successfully, but at the moment not predictably. The use of small systems might aid in the modification of software--to simplify it, reduce it to firmware, and save us from "software collapse".

6. Peripherals will improve in performance at lower cost, but will remain of the same magnetic variety in large measure. Long-range forecasts would indicate the possibility of another peripheral revolution when the laser principles for storage devices are successfully implemented. This must be at least five years off.

7. Networks are "in" today and probably will remain so for quite some time in the broad sense of the word. Computer utilities are a long way off--if ever they materialize. Networks such as those operated by Control Data Corporation and Tymshare are currently in operation; they are doing well and are expected

to grow. They are basically providing services. It is the requirement for services that will determine the extent of competition in resource sharing through networks. Many organizations with their own computer systems today will probably be converted to being customers of service organizations in the future. There seems to be little doubt that eventually these customers will get more service for fewer dollars. Problems yet to be resolved are those of privacy and security. Communication problems that seem to exist today will not be barriers. Communication costs will have to come down to make the remote use of service systems desirable. These services will be provided on an international basis.

COST/INSTRUCTION/SECOND
(DOLLARS)

COST EFFECTIVENESS
OF COMPUTER FAMILIES
(IBM COMPUTERS)

NUMBER OF COMPUTERS

Figure 2

Figure 3