Recent Technical Advances in the Computer Industry and Their Future Impact*

Stuart E. Madnick
Massachusetts Institute of Technology

When one thinks of the tremendous progress in computer development which has occurred during the past 20 years, it is difficult to conceive of a future having as much change. Yet, in this paper Professor Madnick asserts that advances during the next decade will humble those of the last two. He discusses future trends in three major areas: breakthroughs in manufacturing, development of computer system architecture, and efforts to meet user requirements. In addition, key problem areas that may inhibit rapid progress are identified. *Ed.*

After a decade of gestation, computers came to the attention of the general public in the 1960's. It was during this period that many engineers and managers first were exposed to computer systems. Unfortunately, these users have not been able to keep abreast of wide-ranging advances occurring in the 1970's. Only experts in the computer field have been able to understand fully the sweeping changes that are taking place in the industry.

This article surveys these changes. Due to their broad scope, it is not possible to discuss them here in depth. However, it is hoped that highlighting recent developments and projecting their future impact will provide a reference for readers interested in pursuing a particular topic.

The past two decades have seen tremendous advances in computer systems, but those anticipated during the next five to ten years should prove even more significant. These developments can be divided into three categories: technological cost/performance breakthroughs in computer manufacturing; evolution of computer system architecture, both hardware and software; and major steps toward meeting user requirements and capabilities.

Technological Cost/Performance Breakthroughs

Technological developments in computer system hardware can be anticipated to have significant effects in many areas such as cost, performance, size, and reliability. Since there is extensive activity in computer technology, only major trends will be singled out.

1 Microprocessors The development of semiconductor microprocessors is likely to have tremendous impact in many directions. A microprocessor can be defined as

*For those readers who are not familiar with computer terminology, a glossary is provided in the early pages of this paper.

a small computer processor that consists of one or two semiconductor integrated circuits (large-scale integration, LSI).[1] Such a processor's size is measured in inches. A slightly simplified form of this technology has sparked the recent growth of hand-held electronic calculators.

There are no serious technical problems in the development of microprocessors. They typically are based upon standard MOS/LSI (Metal-Oxide-Silicon Large-Scale-Integration) technology. The present problems largely are customer acceptance.[2] The manufacturing cost is very sensitive to production quantity; due to complex engineering and production setup requirements, it can cost only slightly more to manufacture 10,000 microprocessors than 1000. (The joke that the only raw material required is a few shovels of beach sand, though not technically true, is indicative of the situation.) The marketing problem becomes apparent when one recalls that after 25 years there are still less than 200,000 conventional computers in the world.

Characteristics of Microprocessors There are many differences among the various microprocessor approaches. The Intel MCS-4, one of the earliest and simplest, will be described. The MCS-4 is a four-bit-parallel processor with 45 instructions, and can execute almost 100,000 instructions per second. For comparison, a conventional computer, such as IBM's 370/135, has a 16-bit-parallel processor, over 100 instructions, and can execute about 1,000,000 instructions per second. The microprocessor does have advantages in cost. The heart of the MCS-4, the 4004 control and arithmetic unit, sells for 30 dollars in 100-unit lots. The entire four-chip MCS-4 system sells for 48 dollars in 100-unit lots.

The long range characteristics are difficult to predict. Present sales prices probably are artificially low since none of the manufacturers are in full-scale production. However, there are improvements possible in all directions. Eight-bit and 16-bit-parallel processors have been developed. At least one recent reference has predicted that microprocessors with execution rates up to ten million instructions per second will cost about one dollar within 25 years, though such a long range projection is strictly conjecture.[3]

Uses for Microprocessors As noted above, the immediate problem is developing uses for microprocessors. Potential uses can be subdivided into three areas: noncomputer applications, small computers, and new approaches to medium-scale computers.

Noncomputer Applications. In noncomputer applications the microprocessor is integrated into a larger system. This has been done for years with minicomputers (computers selling for ten thousand dollars or less) but will be accelerated by the low-cost microprocessor. The possibilities are limitless.[4] Typical applications include machine-tool control, telephone switching, point-of-sale systems, medical electronics, automotive controls, digital watches, and parking garage fee computers.

---

1 See Lapidus [11].
2 See Rhea [15].
3 See Foster [5].
4 See Lapidus [11] and Foster [5].

Glossary

Bit A binary digit. Most computers operate using the binary (base 2) number system. The decimal number 14 (base 10) becomes 1110 in binary; it is four bits long.

Byte A group of eight bits. Computers usually operate upon groups of bits at a time. The byte 00011110 (base 2) is the same as the decimal number 30 (base 10). Bytes also are called *characters* or *words*.

Instruction A coded number that indicates a step of a computation to be performed. For example, 1175 (base 10) may mean add 175 to a subtotal, and 3010 (base 10) may mean multiply the subtotal by 10. Instructions are coded in binary.

Processor A device that can interpret (decode) instructions and perform the indicated computation. Modern processors can perform millions of instructions per second.

N-Bit-Parallel Processor A processor that can manipulate n bits at a time rather than one bit at a time.

Large-Scale Integration (LSI) An electronic semiconductor technology, based upon the earlier transistor technology, which makes it possible to build devices the size of a postage stamp that contain the equivalent of thousands of transistors.

Main Storage The high-speed, high-cost memory, usually LSI, where current instructions and data are kept for use by the processor. Due to the binary number system, memory capacities often are expressed in units of 1024 bytes (binary kilobytes, K). Thus 4K bytes means 4×1024 bytes.

Intermediate Storage Memory devices that take more time to use (slower) than does main storage, but are much less expensive per byte. Contemporary computer systems consist of many different storage units with different cost/performance characteristics. These devices somewhat arbitrarily are classified as main, intermediate, secondary, or archival memories.

Secondary Storage Devices that are slower than intermediate storage, but much less expensive. These include conventional electromechanical devices such as magnetic tapes, disks, and drums. The major components of a computer system are shown in Figure 1.

Online, Offline A device that connects directly to the computer system is said to be *online,* such as a computer-controlled airline reservation keyboard terminal. A device that separately processes information going to or coming from the computer system, such as a keypunch machine, is said to be *offline.*

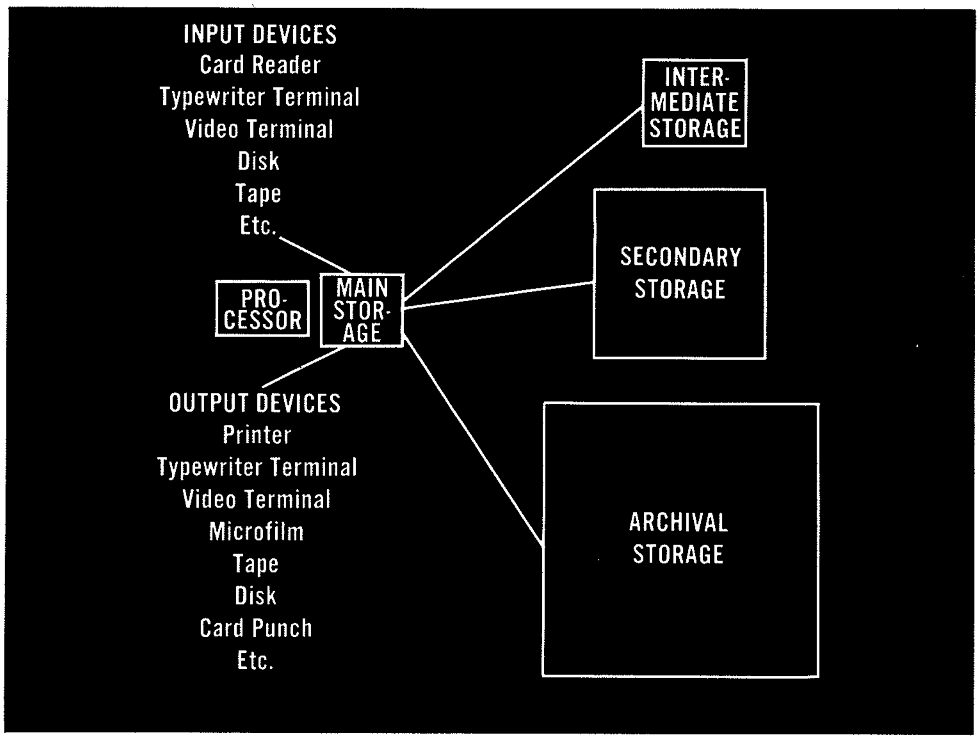

Figure 1 Major Components of a Computer System

Small Computers. There are many applications where a low performance, if low cost, computer is quite adequate. Compared with human skills, a "slow" microprocessor that executes 100,000 arithmetic operations per second is quite impressive. To illustrate the possible effects, it has been reported that Japan's Telephone Company recently cancelled its low-cost touch-tone-input/audio-response-output time-sharing computer business due to the advent of powerful low-cost hand-held calculators. There are many other applications for such personalized and slightly specialized computers both in the home and at work. A considerable amount of entrepreneurial initiative will be necessary to introduce these devices.

In addition to the personalized computer, microprocessors will have impact upon sophisticated remote computer terminals. In fact, much of the current microprocessor activity is jointly sponsored by terminal manufacturers and semiconductor companies. At present, microprocessors are experiencing difficulties in attaining the speeds needed to control high-performance CRT (cathode ray tube) video terminals.

Finally, the larger and more powerful microprocessors could pose a threat to the current minicomputer market. These microprocessors also could evolve into small stand-alone business computers that would process accounts receivable, payroll, etc.

Medium-scale Computers. Medium-scale computers, such as the IBM 370/135, 370/145, and 370/155, represent the major segment of the current computer market. The microprocessors and even their slightly more expensive and powerful cousins, the minicomputers, do not have the processing power or I/O (input/output) flexibility to compete with the medium-scale computer.

However, medium and large-scale contemporary computers are becoming more and more decentralized. In addition to the central processors there are I/O processors (often called *channels*) and special I/O device processors (often called *control units*). Traditionally these processors have all been different, specialized according to their purpose.

As the cost of microprocessors drops and performance improves, one is tempted to use these general-purpose units in place of specialized processors in a medium-scale computer. This trend is already apparent in the recent introduction of the Burroughs B1700 system based upon the Burroughs "D" machine multiprocessor military computer, and the IBM System/370 Model 125 that uses several separate microprocessors internally.

Many factors are likely to accelerate this trend toward multiprocessor and multicomputer configurations. (Technically, "multiprocessor" implies more dependence between the processors than does the term "multicomputer." For simplicity we will use the terms interchangeably.) One intriguing argument is based upon the impact of the rapidly changing state-of-the-art on the development cycle. Typically, the development cycle — the time from initial product conception to full production — is about five years for medium-scale computers. However, for the simpler minicomputer it often is two years, and in some cases less than 12 months. Many important decisions must be made early in the design, such as system performance, circuit types, memory modules, and packaging approaches. The designer of a medium-scale computer must make these decisions four to five years before production. If he chooses to use existing technology, the final product will be several years behind the state-of-the-art. If he extrapolates and projects the availability of future technology, he is likely to make some bad guesses necessitating last minute redesign and resulting inefficiencies. The replacement of the IBM System/370 Models 155 and 165 within two years of initial delivery by the newer Models 158 and 168 is at least partially explainable by this phenomenon. By building medium-scale computers out of general-purpose microprocessors the final design may be slightly less optimal, but this is more than offset by increased flexibility, the ability to use more current state-of-the-art, and cost reductions resulting from larger quantity mass production.

Practical Considerations The dramatic reduction in microprocessor costs must be carefully considered. The direct manufacturing costs of a computer system represent about 25 percent of the sales price or lower.[5] The processor may represent only one-third of the system, the rest being memory and I/O devices. Finally, the processor electronics, excluding cabinetry, power supplies, etc., may represent less than one-third of the processor's cost. Thus, advances discussed above affect less than three percent of a medium-scale computer system's price!

On the other hand, there will be dramatic reductions in the costs of memories and I/O devices. The miniaturized size of the processors will allow considerable reductions in cabinetry, power supplies, etc. Also, the more modular construction impacts other costly areas, such as servicing.

Thus, tremendous advances in processor technology are expected, but the impact upon the end user will be minimal unless there are dramatic changes in these other areas.

2 Main Storage Main storage, also called *main memory* or *core memory,* is a major component of conventional computer systems. It typically represents about one-third of the system's cost. The breakthroughs in this area, both attained and predicted, have received considerable attention.[6]

The semiconductor memory market is extremely competitive in both price and technology. Versions of the Intel 1103, a 1024-bit MOS semiconductor memory circuit, have become standard components and are manufactured by several companies. Due to economies of scale and a very steep learning curve, the cost of such circuits has dropped by a factor of ten in a little over one year. With the eventual commercial maturity of even larger semiconductor memory chips, the cost per bit of memory is likely to drop by another factor of ten in the next few years.

Future computers having larger capacity, smaller size, and less costly main storage can be expected. As a result of the decreased physical size it is likely that more memory will be packaged in cabinets with the processors. As an indication of this trend, the minimum memory size of the IBM 370/135 is larger than the maximum memory size of the IBM 360/30, its predecessor. This increased storage at lower cost will greatly reduce the problems of program development and facilitate the use of sophisticated operating systems.

3 Secondary Storage The use of high-performance direct-access secondary storage was accelerated by the IBM System/360 introduced in 1964. In the past eight

[5] Other costs such as marketing, programming development, research, and profit also are included in the sales price.
[6] See Saxton & Co. [17]; Ayling [1]; Howard [10]; and Thompson, Morton, and Bobeck [18].

years, this trend has continued, and secondary storage has become an area of considerable competitive pressure. Sales of the basic IBM 3330 Disk Storage Unit alone are expected to reach a volume of over one billion dollars. A single 3330 module, storing 100,000,000 bytes, has more than ten times the capacity and double the access speed of its ancestor, the IBM 2311, introduced less than ten years ago.

Current secondary storage devices are based upon rotating magnetic media technology such as magnetic disks, magnetic drums, and magnetic tape strips. This technology probably can be increased by a factor of ten in capacity and two in speed in the next five years. Beyond that point there are at least three limitations to electromechanical equipment:

1 Physical limitations in magnetic media recording capacity and speed are being approached.

2 The cost per bit of storage may decrease as capacity increases, but the unit cost will continue to increase due to the need for costly mechanical motors.

3 The mechanical approach has inherent speed and reliability limitations.

Numerous technologies are being pursued which lead to storage devices that fall between the traditional high-cost high-performance main storage and the lower-cost low-performance electromechanical secondary storage. All of these intermediate storage approaches have been demonstrated successfully in the laboratory, and in some cases in limited production. The most successful should be in full production within five years. These technologies include MOS/LSI shift registers, charge-coupled devices, magnetic bubbles, optical (laser beam) storage, electron beam storage, and bucket brigade devices.[7] Current prices for storage devices and those expected in 1975 are indicated in Table 1.

It is important to note that these intermediate storage devices are both cheaper and faster than the fixed-head secondary devices (e.g., magnetic drums). However, fixed-head storage represents a small fraction of the secondary storage market. It is likely that these new intermediate storage devices will coexist with the faster, though more costly, main storage and the slower, though less costly, moving-head secondary storage. The future of these storage levels will depend upon changes in computer system architecture and applications designed to take advantage of each level's unique characteristics.

4 Archival Storage By standards of a decade ago, today's secondary storage devices have enormous capacity. A single eight-module IBM 3330 Disk Storage Unit has a capacity of 800 million bytes. Assuming that there are about 4000 characters on a dense single-spaced 8½ × 11 inch sheet of paper, 800 million bytes is comparable to 250,000 sheets of paper. Yet, potential information storage requirements greatly exceed this capacity. For example, as part of its antitrust defense, IBM had submitted over 27 million documents as of January 1973.

[7] See Bobeck and Sconil [2]; Fields [4]; Howard [10]; Thompson, Morton, and Bobeck [18]; and Gilder [8].

	Cost (cents per byte)	Typical Unit Capacity (million bytes)	Random Access Time (microseconds)
NOW			
Main Storage	10	.1-1	0.6
Secondary Storage (fixed head)	2	10	5,000
Secondary Storage (moving head)	0.01	100	40,000
1975			
Main Storage	1	.1-2	0.2
Intermediate Storage	.01-.1	1-20	1-1,000
Secondary Storage	.002	200	20,000

Table 1 Current and Expected Prices for Storage Devices

To satisfy these needs there has been considerable activity in the development of archival storage devices capable of storing enormous amounts of information economically. These devices are often called *terabit memories,* since they are designed to hold over one trillion bits of storage (a one trillion bit memory is the equivalent of 40 million 8½ × 11 inch sheets of paper). There are several commercial archival storages already on the market, including Grumman's MASSTAPE, Ampex's Terabit Memory (TBM), Precision Instruments' UNICOM, and IBM's 1360 Photo-Digital Storage (PDS).[8]

These units provide direct access to over a trillion bits of storage with a maximum delay of a few seconds. Typical cost per byte is around .001 cents. In some cases, such as MASSTAPE and TBM, the information is erasable and rewritable — similar to conventional computer storage devices. In many cases, such as UNICOM and PDS, the information is written permanently and not erasable.

Today an archival storage unit has the capacity of 10,000 conventional magnetic tape reels online. The recording media usually is removable and can be stored offline in a fraction of the space required by magnetic tape libraries. The offline storage cost for the recording media itself drops to around .00005 cents per byte. At least one installation was justified by eliminating the cost of purchasing thousands of reels of magnetic tape and providing the necessary storage space.

An 8½ × 11 inch document, assuming 4000 characters of text again, could be stored in computerized form for 4 cents online or .2 cents offline. That is comparable to the price of paper. Thus, it may be cheaper to store information in a comput-

[8] See Dell [3]; Gentile and Lucas [7]; Haines [9]; and Penny, Fink, and Alston-Garnjost [14].

er than on paper! The implications of archival storage units are not yet fully understood, and experiments on their future uses are just beginning. This area could have tremendous impact upon society.

The capacities, speeds, and prices discussed above are already available. In the next five years significant advances can be expected, especially in laser and electron beam approaches used in the UNICOM and PDS units.

5 Other Trends The coming decade will bring many advances in computer-related technology. System reliability will increase due to extensive use of electronic circuits, error checking and correcting techniques, and economical redundancy.

The marriage between computers and communications will intensify. Computers already are being used to control communications in AT&T's Electronic Switching System (ESS). Digital communication, as contrasted with voice, is increasing rapidly. This area is complicated by many factors such as technology, FCC regulations, relatively inexperienced and rapidly growing competition, and AT&T.[9]

Computer System Architecture

In the first section current and anticipated technological advances that will have significant effects upon the ways that computers are used were presented. Another source of major change is in the area of computer system architecture. As a recent conference speaker stated, "After 25 years of growth, the computer industry has reached its infancy." The early computers were primarily high-speed calculators used to generate ballistic trajectories. Now systems are being used for purposes undreamed of 25 years ago. Yet basic computer structures have not changed much over the years. Research during the past decade is about to pay off in new and more effective approaches to computer architecture.

1 Multiprogramming The technique of multiprogramming, the interleaved execution of two or more programs, is standard on most medium and large-scale systems. The procedures presently required to accomplish multiprogramming often are awkward, require a sophisticated operating system, and introduce considerable overhead and performance degradation.

By analyzing the fundamental requirements needed to support multiprogram operation and incorporating these features into the basic computer hardware, operating systems become simpler and more efficient. Rudimentary attempts to accomplish this can be seen in the old Honeywell 800 series and the recent Singer Ten System. Far more significant approaches can be found in the Venus Project at MITRE. Similar experiments exist in the advanced development laboratories of most major computer manufacturers.

---

[9] See Mathison and Walker [13].

Many of these multiprogramming facilities can be introduced without affecting existing user programs. However, to attain even greater effectiveness, especially in a multiple microprocessor environment as explained in the preceding section, it will be necessary to develop new programming styles. IBM's PL/I (Programming Language/One) provides some of the necessary multiprogramming features, and other programming languages are being developed.

2 Microprogramming and Control Hierarchies Early computers were relatively simple, though voluminous, performing additions, subtractions, comparisons, etc. As users developed requirements for advanced mathematical processing such as vector and matrix operations, extensive data base processing, and intricate problem solving, far more sophisticated computers were desired. The microprogramming manufacturing technique makes it feasible and economical to produce such systems.

In the Venus Project mentioned earlier, most of the traditional operating system functions have been incorporated into the basic computer hardware. This approach has also been used to greatly simplify and speed-up the operation of high-level programming languages such as COBOL, FORTRAN and PL/I on the Burroughs 1700 system, and APL on an experimental IBM system. This trend will continue on future systems, thereby providing far more powerful and efficient programming facilities to the user.[10]

3 Virtual Storage and Storage Hierarchies IBM has recently popularized the concept of *virtual storage,* the automatic management and movement of information between main storage and secondary storage. Similar approaches have been used in many earlier systems by other manufacturers such as Burroughs and RCA (now UNIVAC). Virtual storage greatly simplifies the tasks of the programmer — the major cost in application development — and improves system performance.

The effective use of intermediate storage technologies requires an automatically controlled storage hierarchy that provides a virtual storage which encompasses main, intermediate, and secondary storage. Research on this approach is being conducted in the development laboratories as well as universities such as MIT.[11] The current problems probably will be resolved in time to allow the use of intermediate storage devices in storage hierarchies for the next generation of computer systems. The combined effect should further reduce programming costs and increase system efficiency.

4 Communications One of the most significant impacts of information processing systems will be on communications, including time sharing, centralized data bases, and computer networks.

Time sharing, although not meeting the lofty projections of its advocates during the mid-1960's, has had and will continue to have tremendous effects. Most of the ear-

---

[10] See Foster [6].
[11] See Madnick [12].

lier technical problems have been overcome and the reduced system costs make time sharing systems very attractive. Many people who were disenchanted during the expensive time sharing fever of the 1960's are surprised to find that powerful multiple-user online-programming/interactive-problem-solving systems are commercially available for less than $20,000/month from IBM (using the VM/370 operating system on the System/370 Models 135 or 145). More limited time sharing systems, such as those utilizing minicomputers and restricted to the BASIC programming language, are available for a fraction of IBM's price.

The full impact of time sharing systems has been stalled by the lack of entrepreneurial efforts in application areas. Now companies are beginning to emerge which use the time sharing concept to provide useful and convenient facilities directly to the end user. These application areas are as diverse as online advertising media analysis, engineering/manufacturing/production control systems, and sophisticated lens design programs. Many of the larger time sharing services companies already have found that the majority of their revenue is derived from proprietary application program services rather than traditional "raw" computer services.

The identification, development, and marketing of these applications-oriented time sharing services is a serious problem. The successes already attained indicate that these problems can, and probably will, be overcome.

As the complexity of modern-day business increases, it is necessary to rely more heavily upon computer assisted controls. New information handling concepts coupled with the economics of secondary and archival storage devices make centralized data bases feasible. Many of the earlier problems with online real-time total management information systems can be attributed to the naïveté of the users and implementors rather than the basic concepts.

As the need develops for more global optimization of large systems such as the Federal Reserve System, it will become necessary for local computer systems to communicate and exchange information. Advances in this area are being pursued in projects such as the government-funded ARPANET and Michigan's MERIT system.[12] Commercial versions of these systems already are appearing on the market, such as Bolt-Beranek and Newman's Interface Message Processor (IMP), and systems originally developed for ARPANET.

5 Protection The topics of information system protection and security have received considerable attention in the press recently. However, the full implications probably are not apparent to most observers. It is unlikely that anyone would make a serious attempt to steal a company's payroll program or even the customer list — although it is possible. On the other hand, consider a multimillion dollar software development company whose major assets, proprietary application programs, may be represented by a single magnetic tape reel. The lack of effective technical and legal safeguards has been a major obstacle to the growth of the application-

[12] See Roberts and Wessler [16].

oriented time sharing services market. Fortunately, many of these problems have been solved.

When considering the protection of information in a computer system it is useful to divide the problems into three parts:

1 Validation — How to keep bad information out of the system.

2 Integrity — How to prevent information from being destroyed or lost.

3 Security — How to prevent unauthorized access to the information.

It makes little sense to lock an information system in a lead vault guarded by Marines if the information is meaningless or incorrect. A simple example may clarify the point. Recently a New Jersey town, in preparing data for computing the tax rate, incorrectly entered a value of $10,000,000 for one resident's house (misplaced decimal point or sleepy keypuncher). After the tax rate was computed the mayor was pleased to announce that their town had one of the smallest rate increases in the state. Months later, when an irate resident complained of a $10,000,000 property assessment on his eight-room house, the error was uncovered. At last report, the mayor was looking at a sizable budget deficit.

Fortunately, there is considerable activity in these areas. IBM has recently initiated a five-year, 40 million dollar research effort in conjunction with several university, industry, and government investigators.

Systems that are Easier to Use

Advances that are driven primarily by technological innovation in hardware manufacturing and system architecture have been presented above. It is reasonable to ask, "What about the user?" It should be noted that many of the topics in the preceding sections do result in increased user efficiency, effectiveness, and convenience. In addition to the beneficent, altruistic desire to make life more pleasant for the user, there are many important economic reasons that are accelerating the activities of manufacturers and researchers in this area.

However, an underlying problem surfaces rapidly. If the prices of hardware drop as anticipated it will be possible to do the same processing next year for less than this year. One estimator claimed that the current world-wide inventory of processors and memories could be replaced at a cost of less than one billion dollars by 1975. Thus, if there is not continual and massive growth of the market, the industry will stagnate and, in dollar revenues, shrink tremendously.

All the gloomy statements above potentially have applied over the past 25 years. Fortunately, the market always grew much faster than prices dropped, and demand developed for even larger and more powerful systems. When the early ENIAC computer was built, reliable experts predicted that 100 such machines would satisfy the country's computational needs for the rest of the century; needless to say, the market was somewhat larger.

In this paper it is assumed that potential market growth for computerized processing and information systems is enormous. Still, it is possible to identify three major bottlenecks to growth.

Salaries vs. Hardware. In the development and operation of new application areas, it is estimated that close to 70 percent of the cost is tied to humans (including salaries, office space, fringe benefits, etc.) and only about 30 percent to computer hardware. If hardware costs were to drop to zero, there would be relatively little increased incentive to develop applications at a faster pace. The first bottleneck is tied to the cost of people needed to develop new systems.

Maintenance vs. Development. In most mature data processing installations about 80 to 90 percent of the personnel and costs are devoted to the operation and maintenance of existing applications. This leaves only about ten to twenty percent of the budget for the development of new application areas. The second bottleneck is tied to the cost of operating existing systems.

User Sophistication and Education. The two bottlenecks above relate primarily to current, relatively mature users. Another and even larger market is found in the present nonusers. Nonusing firms often are small and unsophisticated with little understanding of computer systems. The third bottleneck is tied to making systems usable by the nonuser.

In this section approaches and techniques that have been developed or are being studied which attack these bottlenecks will be discussed.

1 High-Level Languages and Problem-Oriented Languages High-level languages (HLL) and problem-oriented languages (POL) have been in use for many years. The "machine language" of a conventional computer is awkward and tedious for human use in expressing a problem. Instead, the user expresses his problem in an "English-like" language, such as FORTRAN (FORmula TRANslator) or COBOL (COmmon Business Oriented Language), and the problem is automatically translated into machine language. The net effect is that it is easier and faster for the user to write computer programs.

Past difficulties with HLL's and POL's included high costs associated with building the automatic translators (compilers), difficulties in using some HLL's and POL's, and slow execution of programs written in HLL's and POL's. These problems largely have been overcome by research that has resulted in techniques which produce economical, efficient, and more powerful compilers; the shift in cost dominance from computer hardware to people (i.e., even if the translator were inefficient, a manual translation usually would be much more expensive); and new computer architectures which provide for efficient operation of HLL's and POL's.

2 Generalized Application Packages An examination of typical computer usage reveals tremendous similarities. Everyone seems to want a payroll program, an inventory control program, an accounts receivable program, etc. Why should every company develop its own systems? The vast majority of current programmers are working on projects that already have been done at other companies.

The problem usually is not due to company secrecy. In fact, some companies have found that marketing their programs has brought in considerable extra revenue. The problem, in general, is that no two programs are exactly the same, even when used for similar purposes. For example, the ubiquitous payroll program may differ due to handling of salaried vs. nonsalaried employees, pension plans, etc. Thus, one company's payroll program may be of minimal use to another company.

Although a given application area, such as payroll, may have thousands of variations, there usually is a much smaller number of mutually exclusive options. The numerous combinations and permutations of these basic options result in the tremendous diversity. A generalized application program attacks this problem by providing capability to handle all of the possible options. The user merely specifies which options he needs and what particular values must be used (e.g., state tax rate). IBM has pursued this approach with its Applications Customizer facility used on small-scale System/3 computers.

This approach has made the computer usable and economical for a large market, the small user. It is still necessary to develop these generalized application packages, and they usually are much more costly than any single nongeneralized application program. Thus, only large companies with many customers can justify the initial cost. Furthermore, in areas with even more options such as production scheduling and market analysis, and totally new areas such as medical applications, the cost of developing the generalized program increases and the market size decreases. In the meantime, the growth of entrepreneurial software companies and work done by present computer manufacturers will result in considerable activity in the generalized application package area.

3 Information Handling Actual programming is only part of the cost of developing new systems. In order for them to be used and maintained, tremendous amounts of documentation must be prepared. The use of high-level languages has made some programmers so effective that it takes more people (documentors) to explain what has been done and how to use a system, than to develop it. In such projects the documentation is a serious bottleneck.

Several approaches have been developed to combat this problem. The use of "English-like" high-level languages that are easy to understand reduces the amount of additional maintenance and design documentation that is required. Documentation is required though; the truly self-explanatory program still does not exist.

Just as compilers were developed to help in the programming and translation process, various tools have been developed to aid in the documentation process. Online manuscript processing systems, such as IBM's SCRIPT/370, make it convenient for the programmer to create and update system documentation without the additional costs and delays of requiring separate secretarial services. These tools, although originally developed by computer-people for computer-people, are finding increasing interest and usage in all types of projects that require substantial amounts of documentation.

4 Intelligent Data Bases Computational uses of computers are continually dropping in importance compared with data storage and data processing uses. The decreasing costs coupled with the emergence of archival storage devices are accelerating this trend. During the 1960's many companies made attempts to develop online real-time total integrated management information systems, with disasterous results. In many cases the entire concept and plan for implementation were poorly handled. In other cases existing technology did not meet the necessary requirements.

For example, if a company's personnel data base contains the name of each employee's parents, it should be possible to inquire how many father/son pairs are currently employed. In most conventional systems such a query could not be made unless it had been anticipated or the system were modified to handle it. This example can be extended even further by considering the question of how many grandfather/grandson pairs are currently employed. Note that the grandfather and grandson information is not explicitly stored in the data base, but by using the information on parents and children of each employee the grandfather/grandson pairs can be identified. Systems that are able to make these discoveries automatically are often termed *intelligent* data base systems. There is considerable activity in this area in industry, and in universities such as MIT's Sloan School of Management. These intelligent data bases will enlarge greatly the present information systems market for both large and small users.

5 Intelligent Systems and Automatic Programming A particular generalized application package as described earlier is capable only of handling a specific area. Even within that area it can cope with only a limited number of options. The intelligent data base systems are not specialized to any particular application, but they only handle information-related requirements. They would not be able to produce automatically a complete inventory control system since that involves application-specific background knowledge. By combining the application-specific knowledge of generalized packages with the information-related capabilities of intelligent data base systems, one can hope to develop an intelligent system. Although this is a very difficult objective, it is not quite science fiction. The Advanced Research Projects Agency of the Department of Defense is funding research in related areas, under the heading of "automatic programming," at several universities. One example is MIT's Project MAC Automatic Programming Division. If these projects are successful many of the bottlenecks mentioned at the beginning of this section will be eliminated. Results should materialize within five years.

6 Artificial Intelligence The term *artificial intelligence* often brings to mind robots, sophisticated chess-playing computers, and other far-out sounding concepts. In recent years the field of artificial intelligence, although still very long range, has developed many concrete results. Two significant advances have been *goal-directed programming* and *natural* (English) *language capabilities.*

Goal-directed systems, such as PLANNER and CONNIVER, differ from conventional programming techniques. The user merely expresses, in reasonably precise terms, *what* he wants done rather than *how* he wants it done, which is required for conventional programming languages. This approach makes it easy to build larger and more complex systems, and puts much of the mechanical problem-solving burden on the computer rather than the human. Although these goal-directed systems are still experimental, they have been demonstrated to be effective. In fact, they have provided many of the necessary breakthroughs needed to accomplish the automatic programming described above.

One may visualize true English as the most convenient means of communication with a computer. Although true English may not always be the most desirable approach, recent advances in artificial intelligence indicate that this objective may be attainable, especially in somewhat limited contexts. The Blocks World system, developed at MIT's Artificial Intelligence Laboratory, allows the user to converse with the computer about children's blocks (e.g., which block is on top of the red block, etc.). Researchers and computer manufacturers are exploring ways to extend these techniques to match user needs. An example would be a data base system which can accept real English queries rather than stilted English-like queries.

Concluding Remarks

This report has identified trends in computer system development that are likely to have substantial impact upon the ways that computers are used. The cost and price figures have been provided only as examples. Real costs are highly volatile, difficult to determine, and strongly influenced by production volume considerations. Prices, though available, often are determined by market conditions and have little technical significance. They may be excessively high to cover the start-up production costs or excessively low, even below cost, to develop the market. Since the first version of this article, many of the costs and prices have changed by as much as 25 percent. Although the day-to-day changes may be minor, anyone who must use, evaluate, or coexist with computer systems should plan on yearly or biyearly technology updates. Otherwise, he may be making decisions based upon obsolete or incorrect assumptions.

Much of the research discussed in this paper has not been fully reported in the published literature. As an aid to the interested reader, specific references are made to universities, research groups, and companies active in these fields. The references, though carefully selected, are not necessarily complete nor should they be viewed as a recommendation of the referenced organization.

A key point to consider is that the full potential of these advances is dependent upon dramatic changes in information system usage and structure. In many cases these changes are more under the control of information system users than manufacturers. Thus, an informed and educated user community is an important requirement for future progress in information systems.

References

[1] Ayling, J.K. "Monolithic Main Memory is Taking Off," *1971 IEEE International Convention Digest* (March 1971), pp. 70-71.

[2] Bobeck, A.H., and Sconil, H.E.D. "Magnetic Bubbles," *Scientific American* (June 1971), pp. 78-90.

[3] Dell, H.R. "Design of a High Density Optical Mass Memory System," *Computer Design* (August 1971), pp. 49-53.

[4] Fields, S. "Silicon Disk Memories Beat Drums," *Electronics* (May 24, 1971), pp. 85-86.

[5] Foster, C. "A View of Computer Architecture," *Communications of the ACM* (July 1972), pp. 557-565.

[6] Foster, C. "The Next Three Generations," *Computer* (March 1972), pp. 39-42.

[7] Gentile, R.B., and Lucas, J.R. "The TABLON Mass Storage Network," *Proceedings of the Joint Computer Conference,* Vol. 38 (1971), pp. 345-356.

[8] Gilder, J. "3 Years After Birth, CCDs Head for First Commercial Applications," *Electronic Design* (January 4, 1973), pp. 36-40.

[9] Haines, J.B. "The Evolution of the MASSTAPE System," *1972 IEEE International Convention Digest* (March 1972), pp. 140-141.

[10] Howard, H. "Memories: Modern Day 'Musical Chairs,'" *EDN/EEE* (August 15, 1971), pp. 23-31.

[11] Lapidus, G. "MOS/LSI Launches the Low-Cost Processor," *IEEE Spectrum* (November 1972), pp. 33-40.

[12] Madnick, S. "Storage Hierarchies," *MIT Project MAC Technical Report,* 1973.

[13] Mathison, S., and Walker, P. "Regulatory and Economic Issues in Computer Communications," *Proceedings of the IEEE* (November 1972).

[14] Penny, S., Fink, R., and Alston-Garnjost, M. "Design of a Very Large Storage System," *Proceedings of the Joint Computer Conference,* Vol. 37 (1970), pp. 45-51.

[15] Rhea, J. "Processor on a Chip — Will It Sell?" *Electronic News* (September 25, 1972).

[16] Roberts, L.C., and Wessler, B.D. "Computer Network Development to Achieve Resource Sharing," *Proceedings of the Joint Computer Conference,* Vol. 36 (1970), pp. 543-549.

[17] Saxton, G.A. & Co., Inc. "Computer Memory Technology," *Industry Note No. 12* (March 22, 1971).

[18] Thompson, S., Morton, J., and Bobeck, A. "Memories: Future Storage Techniques," *The Electronic Engineer* (August 1971), pp. 33-39.