

Computers in Science and Technology: Early Indications

Humberto Gerola and Ralph E. Gomory

Computers and their related technology already have become major tools in scientific research and engineering. Almost every aspect of technical activity has been affected in a relatively short time. Computers not only have made existing procedures easier but have led to new ones that were impossible only

matters ranging from funding decisions to how scientists will be able to attack new problems and, even more important, how they will change their methodologies.

The future may indeed be very different from the past in many ways. But looking at the past can help us to avoid a

Summary. Over the past 20 years computer technology has advanced rapidly, especially in the areas of logic and memory. Continued growth at a similar pace can be expected through and probably beyond the 1980's. The technology has already had such major effects on scientific research and engineering that it is of practical importance to try to predict its evolution and uses. It is likely that the trend toward smaller, faster, and cheaper circuits made possible by improved lithographic techniques will continue, resulting in a factor of 10 increase in speed of central processing units and an even greater increase in speed of single-chip microprocessors by the end of the decade. Similar progress is expected in the areas of memory, magnetic storage, printers, and displays. Overall, approximately 20 percent growth annually in the capability of computing systems can be projected. How this continually increasing computing power will affect scientific and engineering activity is more difficult to predict, but some patterns are emerging. Observations of technical personnel at the IBM research laboratory at Yorktown, New York, where the average user has access to a large amount of computing capability and to a worldwide computer network, indicate that workers in different areas have reacted to computer technology in different ways. Whereas engineers have used computing power, displays, and the ability to communicate or share information more or less equally, management has used communication most and scientists have put the greatest value on computing power and displays.

two decades ago. This technology continues to evolve at such a rapid pace that it is important to try to predict its evolution and assess its impact. In fact, with such a pervasive and rapidly changing technical environment, as with any other environment, trying to chart a course by looking forward becomes a practical necessity. It affects the way we think about

narrow focus on a single element of change. This is particularly important when dealing with the complex web of interactions between technical and sociological factors that constitute the scientific and engineering enterprise. The history of technology is full of examples of initial misperceptions about the potential uses and possible consequences of

new technologies. Over and over, technologies that were expected to provide new possibilities in a narrow field ended by having unintended and unpredicted effects in other areas. In the case of computers we have had more than two decades of experience with the technology and its uses, and some specific patterns are already emerging. This article does not pretend to be a complete discussion of the issue of the impact of computers in science—that is too big and diffuse an issue. Instead, we use a few examples from our experience to illustrate the patterns that we are observing.

Evolution of Technology

The most important single factor that has affected and will continue to affect the use of computers in science is the unrelenting, steady advance of technology, especially in the areas of logic and memory. This is due primarily to the fact that every year, through improved lithographic techniques, circuits are made smaller and therefore faster, while at the same time cheaper.

The general-purpose uniprocessor is a convenient basic unit of computing. Figure 1 shows, as a function of time, the number of MIPS, or millions of ordinary instructions per second, that can be executed by a single central processing unit (CPU) of the high-end type. The curve shows steady growth since 1970. In a semilogarithmic plot, it is a fairly straight line, indicating exponential growth. There are good technical reasons for projecting that straight line at the same slope to the year 1990. The processor of today, which can execute approximately 10 MIPS, will evolve through the constant improvement of circuitry, mostly through its miniaturization, toward 100 MIPS.

At about the end of this decade some problems may appear in the form of physical limiting features. By that time the processor will have to be kept very

Humberto Gerola is a manager at the IBM Research Laboratory in San Jose, California 95193, and Ralph E. Gomory is vice president and director of research at IBM, Yorktown Heights, New York 10598.




Fig. 1. Number of MIPS that can be executed by a single central processing unit of the high-end type, projected to 1990.

small; otherwise the time required for a signal to go from one part of the machine to another will start to slow it down. The smaller the machine, the harder it is to remove the heat dissipated by the fast switching devices. Nevertheless, there are reasons to believe that the problems encountered will be solved and that the high-end uniprocessor will improve throughout the decade and probably beyond it.

Much larger increases in performance can be obtained if, instead of a general processor, machines with specialized architectures and organization are used. The increase in speed and above all the decrease in price of the circuitry make it easier to develop more powerful machines for more restricted sets of problems. The degree of improvement increases with the degree of specialization one is willing to accept.

Advances in the miniaturization of circuits will also affect the area of non-high-end machines, or single-chip microprocessors, and here progress will be even faster. Figure 2 shows the number of transistors on single-chip microprocessors. The number of MIPS that




Fig. 2. Evolution of the number of transistors on a CPU chip.

can be extracted from a one-chip processor increases more rapidly because the individual circuits become faster and because more circuits can be crowded onto a single chip—that is, there are two sources of progress. By the end of the decade we should see 10 MIPS from a single general-purpose chip, compared to about 1/5 MIPS in common use today.

Similarly, the progress made in the area of memory is depicted in Fig. 3. By the end of the decade a dynamic memory or ordinary read/write memory of 1 million or 2 million bits to a chip is foreseeable. Again, there has been more or less steady growth since the early 1970's.

It is somewhat more surprising to see the same trend in the area of magnetic storage, which includes disk recording devices (Fig. 4). Unlike semiconductor technology, which has advanced through improvements in lithographic and processing techniques, storage seems to involve an antique and mechanical technology. Disks are basically like phonograph records. They have an arm, which has on it a sensitive head. The head flies over the surface of the rapidly rotating disk and senses changes in the magnetic field caused by magnetic bits recorded on the surface. In present disks the head flies at a height equal to a fraction of a wavelength of light. Progress has continued by decreasing the height of the head over the disk so that it can detect ever smaller bits recorded on the surface. Detailed technical studies of this area tend to conclude that the straight-line progress of the past will continue through the end of the decade.

In communication there is no simple index of progress but no reason to anticipate a bottleneck. We have emerging technologies, such as fiber optics, with great bandwidth potential, and satellites for long-distance communication. Both of these technologies will lower the cost of broad bandwidth communication. In addition, optical fibers may be used as a component not only for what one normally thinks of as communication but also for linking one part of the computer to another, for example, linking the main memory to disks.

In discussing computing systems one must consider not only the logic, memory, and storage but also the printers and displays. Our best projections for these are also for continued rapid progress.

Displays can be roughly characterized in a two-dimensional fashion, as shown in Fig. 5, where the horizontal axis represents increasing resolution or number of points in the image, and the vertical axis represents the degree of interaction with the display. At present, a special




Fig. 3. Evolution of dynamic RAM (random-access memory) density.

dedicated processor, usually a single chip, enables the display unit itself to change the data that are displayed on the screen. The power of the processor determines the degree of interaction with the screen. Producing a simple letter takes very little computing power dedicated to the display, whereas rotating a three-dimensional object on the screen calls for a powerful processor. Hence, to some extent progress in displays is a by-product of the rapidly evolving processor technology and can ultimately be credited to the rapid evolution of semiconductors and smaller and faster circuits.

Increased resolution is more dependent on ingenuity, but so far we have been able to produce technologies that enable us to display more and more characters, and the evolution of displays for more purposes seems fairly straightforward.

Printers are an even more peculiar area. Printer technology seems to be characterized by the ability to evolve a specialized device for almost every small niche in the computer market. In Fig. 6




Fig. 4. Evolution of areal density of magnetic storage.

different types of printers are represented in terms of resolution on the horizontal axis and of speed on the vertical axis. The gap in Fig. 6 is a significant area for development. It represents the relatively high-resolution (letter-quality) printers of moderate speed that will be required to go with the intelligent terminals of the future. In fact, the gap exists in Fig. 6 only because there are many new, and in many cases proprietary, contenders for this area.

The message in all this is very simple. As we review the components of the computing systems that will be available to scientists in the next 10 years, we have in some sense a rather dull story: approximately 20 percent growth from year to year in the capability of the system. This predicted increase seems almost certain to have a profound impact.

Impact of Technology

Technology evolution is extremely hard to predict. Anyone who has done such predicting and had to live with the results can attest to that. But trying to predict the effects of technological change is even more difficult. As an example, consider a question that has been addressed by many others, namely, how did the introduction of the telephone affect urban congestion? Looking forward, it would be possible to argue either way—that the telephone would promote urban congestion or that it would reduce it. It might promote urban congestion because the business office would no longer have to be right next to the manufacturing plant; it could be moved into the city. Once the office location was uncoupled from the manufacturing location the office would migrate somewhere else, most probably to a central location. The opposite argument is that once you have the telephone your office no longer needs to be centrally located because you can make up for not being proximate by being in touch over the wires. In other words, in looking forward there would be a problem in trying to predict the impact of the telephone.

The surprising thing is that looking backward is just as bad. A very interesting series of articles on the subject (1) concludes that, after looking at the history of the telephone, its effect is difficult to discern. Not that there was not an effect, but what it was is far from obvious.

Attempts to forecast the impact of technology in general, and of computer

technology in particular, should therefore be made with a great deal of modesty. The tendency to predicate cultural evolution on technological evolution alone must be avoided. All that we really have to go on is what actually has been happening. Figures 1 through 6 not only show rapid growth into the future, but also indicate that we have experienced more than a decade of that same rapid growth. It may be time to ask whether there are already patterns that indicate how this evolution interacts with other factors that determine the evolution of scientific activity.

The most useful thing to do at this stage is to concentrate on a few specific examples. We will describe one of the large research laboratories of the IBM company and the effects that we have been able to see on our population. The laboratory, located in Yorktown, New York, is part of our Research Division, which also has laboratories in San Jose, California, and Zurich, Switzerland. It conducts basic research and engineering and applied research. In addition to these research laboratories, there are some 26 product development laboratories in IBM worldwide.


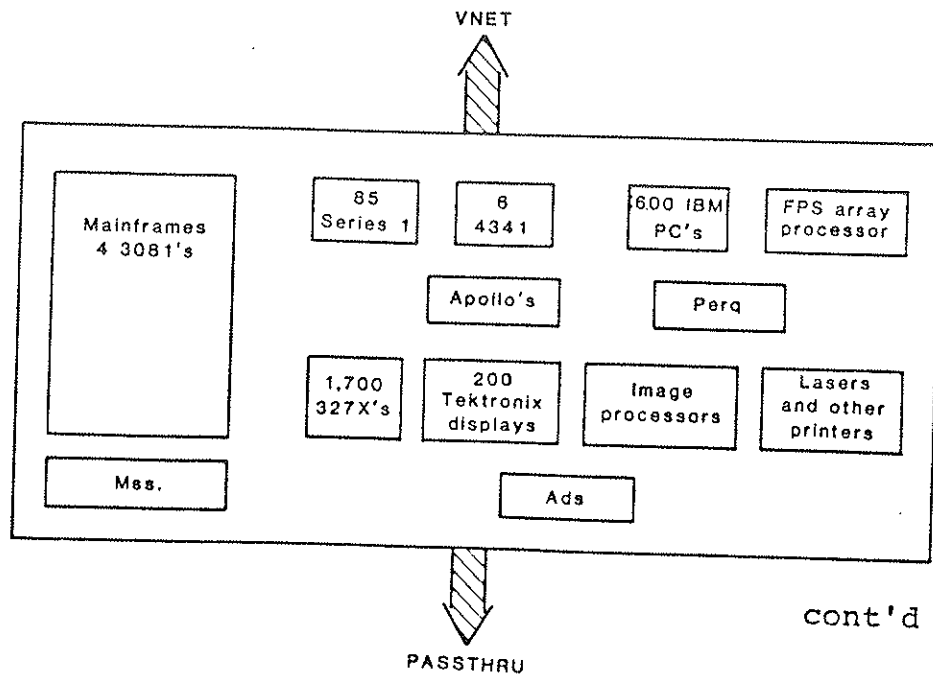


Fig. 5 (left). Two-dimensional characterization of displays. Abbreviations: CAD, computer-aided design; WP, word processing; DP, data processing. Fig. 6 (right). Two-dimensional characterization of printers.



cont'd

Fig. 7. Computing power at the IBM Yorktown laboratory.

The Yorktown Laboratory

The population of the Yorktown laboratory includes some 1200 technical people—computer scientists, electrical engineers, physicists, mathematicians, chemists, and materials scientists. The computing power to which they have access is shown in Fig. 7. There are four large mainframes at the high end of the IBM processing line. Distributed around the building are almost 2000 display terminals connected to the large processors. At present there are some 600 "intelligent" terminals, mostly IBM personal computers (PC's) and a few higher capability Apollo's and Perq's. In addition, there are 85 Series 1 (S1) minicomputers, largely dedicated to laboratory automation, six smaller IBM processors, and a host of specialized memory-type displays, image-processing systems, and different kinds of printers. Attached to one of the mainframes there is a Floating Point Systems array processor.

To give some feeling for the evolution of this computing facility, we may add that when one of us (R.E.G.) joined IBM's Research Division in 1958, the total computing power of the division

was equivalent to about one and a half times that of a single IBM PC.

Constituting an important element of this complex are the communications facilities. The large CPU's, and of course the smaller ones hooked to them and the terminals, are connected through networks to other IBM facilities. In particular, through VNET (Fig. 7), they are tied into a worldwide network that links the research and development laboratories and many other offices. At present, this network links about 1000 CPU's, and it is growing at the rate of 1 CPU a day. It connects about 100,000 people at 140 geographically distinct locations by lines that vary in speed from 4.8 to 56 kilobits per second. This network is in the process of being connected to external non-IBM networks. The services include electronic mail, file forwarding, and a directory of users. The network PASSTHRU, which is smaller, allows interactive access to the CPU's in its nodes. Complementing the communications facilities at Yorktown is the Audio Distribution System, a voice store and forward system using telephones and supported by an S1 minicomputer. The computing capability provided to the av-

erage user at Yorktown is roughly equivalent to a 2-MIPS machine with 2 to 4 megabytes of memory and a virtually open-ended ability to store files. This total communicating and computing effort is very significant to us. In fact, we spend about 15 percent of our total research budget on this kind of facility.

In discussing the effects of these facilities on our basic research and engineering and applied research activities, we will distinguish three basic elements of computing systems. They are computing power, displaying power, and the ability to communicate or share information.

Basic Research

In basic research, computing power is what matters most. Scientists like displays and use them in novel ways. Sharing does not today figure in a significant way, but MIPS have been important in science since the beginning of computer technology. As an example, we will look at one specific case from solid-state physics and discuss some of the reasons why this is so.

It might seem remarkable that there are still questions about the atomic structure of silicon, one of the materials most commonly used in technology. But it is still not clear how atoms position themselves on the surface of silicon and other materials. The geometric relationship of atoms in the surface can be expected to be different from that in the bulk. In either case, the atoms are in equilibrium under the forces of neighboring atoms, but in the surface half of the neighbors have been removed and therefore the positions of the atoms must be different. Until very recently, it was not possible to resolve this experimentally.

A commonly accepted idea about semiconductor surfaces was that alternate atoms were displaced toward and away from the surface. It was thought that such buckling would lower the energy of the arrangement, which would therefore be a plausible one for the atoms to settle down in. A series of massive calculations was undertaken. Testing a single detailed model took one to several weeks of running for several hours a night on our high-end machines. The results of the computations completely contradicted the accepted notion; the buckling, in fact, raised the energy of the configuration (2). These results were widely accepted and led to many researchers abandoning work based on the buckling mechanism. Subsequent large-scale calculations suggested a new candidate mechanism, called pi-bonding, by




Fig. 8 (left). Spiral galaxy NGC 1232. Fig. 9 (right). Model of formation of arms in a spiral galaxy. The black dot represents an exploding star.





Fig. 10. Computer simulation of spiral galaxy formation.





Fig. 11 (left). Artificial mountain valley band on a fractal mathematical model. Fig. 12 (right). Artificial earthrise over a lunar landscape based on a fractal model.



which surfaces can lower their energy. There are two specific points that we wish to make here. First, the computations described above are not a peripheral aspect of surface science but are at the heart of the subject. Second, the results based on these computations have been accepted.

Numerical computer calculations were essential in establishing the validity of the new ideas and were crucial to their acceptance, which has been very rapid. They had a large influence on experimental work because they were performed on a model of very well established credibility. The calculations consisted of the self-consistent solution of fundamental equations for the electron density and total energy. They involved no approximations beyond those of the general theoretical framework and hence they benefited from the credibility created by successful applications of the same theoretical framework to a wide variety of other similar problems. The computer allowed numerical experiments to play the same logical role in the scientific dialectic played by real laboratory experiments. The advances that we discussed earlier which might give us a factor of 10 improvement in MIPS over the next decade would make possible significantly more work along these lines. What we have tried to illustrate is that MIPS matter and that more MIPS will matter even more.

The end product of large-scale computations is normally large masses of data. It is hard to rationalize or make sense of the results from a printout of streams of numbers. This is why displays are the second most important element of inter-

est for scientists. A graph or a map is often an easier way to understand a great deal of data. Something that is a little different, and probably a harbinger of things to come, is the use of displays not only as a way to summarize vast quantities of numerical data but directly as a way to assess the correctness of a piece of scientific work—not through the interpretation of numerical data but through the direct appreciation of the form on the screen. Two examples worked out at our laboratory will make this clearer.

An outstanding problem in astrophysics is that of explaining the shape and evolution of galaxies. Figure 8 shows a galaxy of the familiar spiral type. The entire disk is rotating, the inner parts with greater angular speed than the outer ones. One mechanism proposed for the formation of the spiral arm delineated by bright young stars is illustrated in Fig. 9 (3). The mechanism is driven by the explosion of very bright stars, which tend to blow up the gas surrounding them and compress it in thin shells at a distance away. These compressed shells of gas may provide the conditions needed for the creation of new stars. The black dot in Fig. 9 represents an exploding star, and in the areas surrounding it there is a certain probability of its explosion setting off the birth of other stars. This process can then be repeated as in a chain reaction.

It is possible to follow this kind of mechanism by simulating it on a computer. That has been done, with the results shown in Fig. 10. With the proper choice of parameters, the appearance of the model galaxies is very similar to that of real galaxies. A first test of the plausibili-

ty of the proposed mechanism is thus the direct visual appreciation of the formation of the spiral arms. Given the strong nonlinearity of such a mechanism, it would have been much more difficult to show, in an analytical or numerical way, that these explosions would generate such distinct spiral arms.

A similar use of displays is in the area of fractals. It has always been difficult to characterize and generate shapes. According to a view proposed by Mandelbrot (4), many familiar natural shapes such as that of a coastline or a river are not, in fact, one-dimensional. Certainly, they are also not two-dimensional. Mathematicians have long been familiar with the idea that between one- and two-dimensional objects there are objects of fractional dimension. However, that many natural objects might best be regarded as having fractional dimension is a new idea. If one tries to measure the length of a coastline and does so in greater and greater detail, following every in and out with greater and greater fineness, the length tends to grow without bound. It turns out that it is useful to treat such natural shapes as objects of fractional dimension.

However, in constructing models of coastline or of other natural objects such as islands or mountains, displays are important. Again, the most direct test of the validity of the model is not a sequence of numerical parameters (because it is difficult to characterize a natural scene by a few key parameters) but the way it looks. Figure 11 shows an artificial mountain valley constructed by use of a fractal mathematical model. By varying a few key parameters in the

model, different types of mountains can be made to appear. The realistic look of this scene is a direct test of this kind of theory. Another and even more spectacular example is the artificial earthrise over a lunar landscape shown in Fig. 12. Again, the mathematical model has been verified by directly viewing the resulting picture on the screen.

The third and least important dimension of computing systems, as far as scientists are concerned, seems to be that of sharing. In spite of the excellent facilities at our Yorktown laboratory for sending and receiving messages, communication by computer is not an important aspect of our scientific work or, as far as we can tell, that of others. Between 1969 and 1982 scientists from the three locations of the IBM Research Division published almost 10,000 reports in the open literature. Of these, only 73 were coauthored by people from different locations, and about half of those were collaborations resulting from long-term visits in which one author was residing at the other location. There has been not even one paper coauthored from the three sites, although they are well linked electronically. Almost all the papers with more than one author are the product of collaboration in face-to-face situations. These papers include all areas: pure science, applied science, computer science, and so on. Therefore, the numbers quoted above are a strong upper limit to the number of collaborations in science and an even stronger upper limit to the number of collaborations in science due to the use of the network.

Over the period cited, while communications and other facilities were being constantly improved, there was no apparent trend toward or improvement in collaboration between scientists at different sites.

A bibliographic search carried out by the Institute for Scientific Information has shown that the same pattern exists in another organization which, like ours, has a large number of scientists doing research in an industrial setting at several locations. The existence of an electronic linkage, powerful though it may seem, has not, in fact, altered the traditional methods of scientific work. When we examine engineering activity, we will see that there is much greater intersite collaboration than in basic research. Anyone is free to speculate on this observation. In our experience it is a fact, however, that electronic communication, even when given away free, has not yet altered the fundamental way in which scientific work has been done. Face-to-face communication, so far, appears to be essential to scientific collaboration. Perhaps we should not be surprised when we consider that this largely individualistic pursuit has survived, in more or less unchanged form, the assault of the postal service, the airplane, the telephone, and the automobile. Perhaps even more remarkably, it has survived the change of scale of science itself, from an activity carried out by a very small number of people to one involving thousands and thousands of researchers. It appears that it may well survive electronic communication.

Engineering and Applied Research

In engineering, the situation is quite different. Whereas for scientists it is computing power first, displays second, and communications a poor third, our experience indicates that in the engineering disciplines, with their much greater tradition of shared project work, all three elements of computing matter in a much more balanced fashion.

The requirement for computing power in engineering is obvious. There has always been a great demand for simulations of complex designs and processes, including computers themselves. We mentioned before that very high computational performance can be obtained by building specialized machines tailored to a specific problem. An example in our laboratory is a special processor built to simulate, with a high degree of parallelism, the logic of a computer.

In the area of displays, the use of computer-aided design is very well known. There is also an increasing use of three-dimensional modeling in engineering. Figure 13 illustrates a three-dimensional modeling application. It shows a high-frequency power supply designed for use in a large computing system. Figure 14 is an exploded view of the model of the previous figure. The point here is that in the engineering world, in addition to the familiar two-dimensional objects, we are starting to work directly with the three-dimensional pictures.

Finally, the third aspect of computing, sharing, has already become an essential element in engineering. One example





Fig. 13 (left) Computer-aided design of a high-frequency power supply for use in a large computing system.

Fig. 14 (right) Exploded view of

from our experience is in the area of the design of chips. Chips are made by creating masks for the lithographic process, which are essentially pictures of various layers in the silicon. They are tremendously complex, as there can be more than 100,000 transistors on a single chip. The data that go into each mask are stored in a computer, and this common database is accessed by the large number of engineers, who contribute individually to forming the mask. This kind of sharing is a commonplace of engineering today and is true of other aspects of chip design.

In software, collaboration of this sort is also routine. A compiler development involved the sharing of work between a California laboratory, the Yorktown laboratory, and an outside software company, with versions of the program transmitted back and forth continually between the three locations through the network. Various versions of a program under development are centrally stored, and the computer scientists working on it have access to it to update the individual versions and make changes. Software development today is often dependent on this kind of sharing.

Management

In the industrial research community there is a third class of people associated with scientific activity, and that is management. These are the people, mostly scientists and engineers themselves, who are responsible for the execution and coordination of the large variety of projects. For management in general, not only scientific management, the emphasis is not on MIPS or displays but on sharing.

In order to keep up with what is going on in a large research laboratory, mail systems, both text and audio, are extremely useful. One advantage is that they desynchronize communication. When you have an idea or want to know something, you can send your message off and it does not matter whether the people you send it to are there. When they come in or are available, they can find your message and reflect on it and reply. Another advantage is that of addressing a large number of recipients simultaneously. After registering your message only once, you can send it to any of those on a given list of people. These tools are very important to us already, and we expect that they will become widely used and will be major communication tools for management.

Discussion

To summarize, among the three populations that we have had experience with, for scientists MIPS come first; for engineers MIPS, displays, and sharing all play a role; and for management, communications is clearly more important. Are these patterns indicative of fundamental cultural differences, or simply transient reactions to a rapidly changing environment?

All aspects of computer technology will continue to evolve at a rapid pace. Figure 15 shows schematically our view of the computing system of the future. It is a complex of powerful engines connected in a network by good communications facilities. There is a central data-processing (DP) complex in which the 100-MIPS machines described earlier are located; hooked up to them are specialized processors, designed especially for engineering and scientific use. Scattered around are smaller processors, to which intelligent processors based on single microprocessor chips with a power of perhaps 10 MIPS are attached. Local area networks are hooked through a gateway and through communications to other systems, including the large one. Intelligent workstations (IWS) are connected to the network through a private branch exchange (PBX) and also to a number of intermediate machines that

play a role as departmental processors or communicate directly with each other through a peer-coupled system. In addition, the network will transmit not only printed messages but also images and voices. Everything we know how to do today will still be done, but with a factor of 10 improvement in power. In addition, there are some things that are possible, though harder to predict, such as symbolic rather than numeric calculation and novel logic-based types of software such as expert systems. These requirements may lead to machines specialized for these needs.

More MIPS will mean, as in the solid-state example, that more problems become tractable. More displays, higher resolutions, and greater interactivity will mean that novel ways of using the displays, such as three-dimensional and other more complex techniques, will become more significant. Increased sharing should lead to better management and the use of project-sharing techniques worldwide.

These are the simple straight-line projections for the evolution of the technology. Its impact on various research activities is in the much more difficult realm of qualitative projections.

Will another factor of 10 cause scientists to cooperate and communicate through computer networks as engineers already do? It may be that engineers and




Fig. 15. View of a computing system of the future.

Computers...

computer scientists are more aware of the potential of the present systems and are willing to put more effort into using them, while pure scientists, for whom the computer is another tool, have a lower level of pain. If this is the case, it may be only a matter of time before everybody operates in the same mode. However, one can make the following observation: scientists, either in the laboratory or in computing, have shown that they will push their systems or tools to the limit in order to get to the results. In computing they are willing to learn to program in machine language if that gives the performance they need for a specific problem. We are now seeing physicists developing and building their own special-purpose calculating machines at a great cost in time and effort. In the laboratory it is common for scientists to take commercial instruments apart and rebuild them to improve per-

formance, again at a great cost in time and effort.

In our laboratories, pure and applied scientists have access to the same facilities, but their patterns of collaboration are very different. It may well be that we are dealing here with subtle but strong cultural factors. It is easy to develop theories of why this is so, but it is difficult to decide one way or the other. This is a fascinating and important subject but more work, and perhaps more experience, is required to understand the reasons. Similar questions arise in connection with other fields that have proved intractable. For example, will education, that crude process in the classroom that has withstood every technical assault for the past 2000 or 3000 years, finally crumble before the impact of electronic progress? Some people think so and have projected that the interaction of computers with instruction

will do it, but still we do not know. Will the availability of terminals in the home, the ability to program at home, and the ability to interact with others over wires, over glass, or possibly through satellites fundamentally change the working patterns of people? That is certainly possible, and again we do not know. Our inability to understand and predict the qualitative effects of computer technology is great. But even the straight-line projection, from what we have experienced to what we can reasonably expect to be the impact on science, is impressive.

References and Notes

1. I. de Sola Pool, Ed. *The Social Impact of the Telephone* (MIT Press, Cambridge, Mass., 1981).
2. K. Pandey, *Phys. Rev. Lett.* **47**, 1913 (1981); *ibid.* **49**, 223 (1982).
3. H. Gerola and P. E. Seiden, *Astrophys. J.* **223**, 129 (1978).
4. B. Mandelbrot, *The Fractal Geometry of Nature* (Freeman, San Francisco, 1982).