THE COMPUTING REVOLUTION AND THE PHYSICS COMMUNITY

By any measure, society— US society in particular is now in some phase of what is being called the information revolution. This "third wave," a term introduced by Alvin Toffler, follows the industrial revolution of the 18th century and the agricultural revolution of 10 000 years ago, and is resulting in phenomenal social, cultural

phenomenal social, cultural and economic transformations. It is changing the ways in which people spend their time at work and at play, and the ways in which they live their lives in general. Inevitably, it will lead to major geopolitical upheavals. This revolution, as were the two previous ones, is technology

based, with the seeds being planted over a long period. The seeds date back to the dawn of human development and tools such as the triangle, compass and abacus. More modern contributions include Napier's bones (1617) and the slide rule that evolved from it, Blaise Pascal's adding machine (1642), Charles Babbage's difference engine (1822) and Vannevar Bush's differential analyzer (1931). During World War II the need for powerful computational tools came from two sources. One was the designing of atomic weapons at Los Alamos (now Los Alamos National Laboratory). The other was improving methods of cryptography, with the need to decipher the Enigma-encoded messages used by the Germans and to break the Purple Code used by the Japanese.

A good choice for the beginning of the information revolution might be the commissioning in 1946 of the Electronic Numerical Integrator and Computer, the first fully electronic computer.

The physics and mathematics communities have played a major role in the long development period of this revolution, going back to ancient times, as well as during the 50 years that have elapsed since the ENIAC.

The formative years (1946–55)

During the years immediately following the war, a number of computers were designed and built as new ideas and technologies were introduced, mostly under government sponsorship. The construction of the machines took place on university campuses, in government laboratories and at commercial firms beginning to take an interest in a nascent industry.

ALFRED BRENNER is deputy director of the computer and software engineering division of the Institute for Defense Analyses in Alexandria, Virginia. He was the head of computing at Fermilab from 1970 to 1984.

The information revolution that the ENIAC heralded 50 years ago has profoundly changed the ways in which physicists—along with society in general—work and interact.

Alfred F. Brenner

The ENIAC was built at the University of Pennsylvania by physicist John Mauchly and electrical engineer J. Presper Eckert Jr for the Ballistic Research Laboratory of the US Army Ordnance Corps. Its design was based on concepts used in the Mark I computer built by IBM in 1944 for Howard Aiken, an applied mathema-

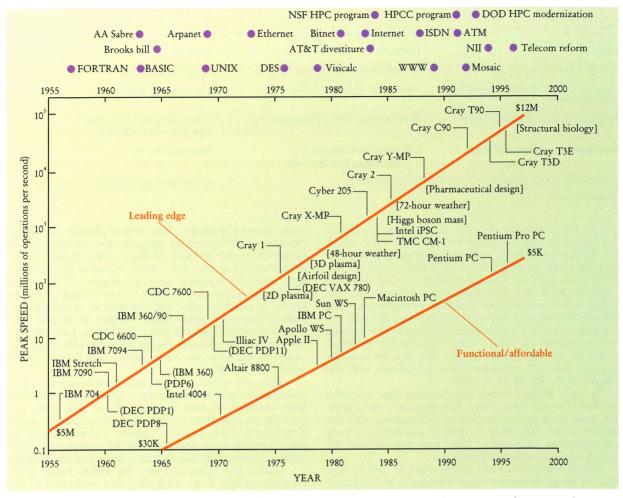
tician at Harvard University, and on the work of John Atanasoff, a physicist at Iowa State College (now Iowa State University) who had been working since the mid-1930s on digital methods of calculation and the use of vacuum-tube circuits for computing.² (See Allan Mackintosh, "The First Electronic Computer," PHYSICS TODAY, March 1987, page 25, and Robert Seidel, "From Mars to Minerva: The Origins of Scientific Computing in the AEC Labs," on page 33 of this issue.)

The idea of the stored program was developed during the design stage of the Electronic Discrete Variable Automatic Computer (EDVAC), which was built in 1950 at the University of Pennsylvania for the Ballistic Research Laboratory. The ENIAC had demonstrated a need to store data, and the concept of also storing the program gave rise to enormous flexibility in program execution. Two types of storage devices were under development, making it possible to build such machines. They were the Williams tube, using electrostatic storage on the faceplate of a cathode-ray tube, and acoustic delay lines in mercury columns.

Other computers built during this formative period³ included the Standards Electronic Automatic Computer, a machine similar to the EDVAC and built in 1950 at the National Bureau of Standards for the Air Force. In 1952, John von Neumann built the Institute for Advanced Study Computer, initially sponsored by the Army Ordnance Corps and later by the Office of Naval Research, the Air Force and the Atomic Energy Commission. Similar machines were built at several AEC laboratories, including Los Alamos National Laboratory, Argonne National Laboratory, and Oak Ridge National Laboratory. Other university machines were the Electronic Delay Storage Automatic Computer (EDSAC) at Cambridge University in 1949, and the Whirlwind in 1951 at MIT.

Industry was also active. In 1950, Engineering Research Associates built the Atlas, the first US computer to use a magnetic drum memory, for the National Security Agency.⁴ In 1951, Remington Rand built the UNIVAC 1 for the US Bureau of the Census. And in 1954, IBM built its first all-electronic stored-program computer, the IBM 701.

By the mid-1950s, computer pioneers had demon-



GROWTH OF COMPUTER TECHNOLOGY SINCE 1955, showing advances in average commercial computer performance and milestone events. Those milestones shown above the upper, or leading-edge, curve are commercial supercomputer products, and those below the curve are other important related events. Milestones in parentheses are not leading edge. Note that individual milestones do not necessarily fall on the curve, although they are depicted that way. The lower curve shows milestone dates of important processors at the affordable level, first for experiments (for example, the PDP8 in 1966) and later for desktop computers (for example the Pentium PC in 1994). Shown in brackets are computational problems that are solvable in reasonable times at the indicated level of computer performance. Approximate mid-level system prices are shown in dollars of the time. At the top of the figure are milestone dates, marked by dots, for events of related technology, software, legislation and so on. In almost all cases, dates and performance levels are approximate. FIGURE 1

strated that it was possible to build quite large computing machines containing thousands of vacuum tubes. Also, the architectural concept of the stored program had been introduced. But most important of all, in 1947, at AT&T's Bell Telephone Laboratories, John Bardeen, Walter Brattain and William Shockley had invented the transistor, which was to prove to be a critical event in the dawning of the information age.

Start of the modern computer era (1956–65)

In 1956, IBM's high-end product was the IBM 704 computer, whose design had evolved from that of the IBM 701. This computer had a ferrite core memory, which had recently been invented by Jay W. Forrester. Engineering Research Associates' Atlas II, installed at the National Security Agency in 1954, was the first computer to use a ferrite core memory. The IBM 704 memory stack contained 32 768 36-bit words and cost over \$1 million—about \$1 per bit. Magnetic drums served as a secondary storage system. Like most computing machines built up

to that time, there was no operating system. There was, however, a symbolic assembler program (SAP), which made it possible to program the machine.

By that time, however, it had become clear that a language programming tool was required if these ever more powerful computers were to be used effectively. For the IBM 704 computer, IBM provided a higher-order-language formula translator program, FORTRAN, developed under the leadership of John W. Backus and a dozen collaborators. Similar efforts were ongoing elsewhere in the US and UK, but it was the delivery of a successful FORTRAN compiler along with a primitive operating system in 1957 that was another significant event in the maturation of the emerging computer industry. It was now possible for scientists to program their problems with relative ease.

In 1960, IBM replaced the IBM 704 and its successor, the IBM 709, with a solid-state version of the same architecture, the IBM 7090. The vacuum-tube—based electronic computer was now obsolete. Magnetic core re-

NSF High-Performance Computing Centers (1985-present)

Center	Location	Grantee
Cornell Theory Center	Cornell University; Ithaca, N.Y.	Cornell University
National Center for Supercomputing Applications	University of Illinois; Urbana, Ill.	University of Illinois
Pittsburgh Supercomputing Center	Westinghouse Corp.; Pittsburgh, Pa.	Carnegie Mellon University/ University of Pittsburgh
San Diego Supercomputer Center	University of California, San Diego; La Jolla, Calif.	General Atomics
John von Neumann Center*	Princeton, N. J.	Consortium for Scientific Computing

*No longer funded

mained the technology for the random-access memory. The IBM 7090 and its successor, the IBM 7094, which was introduced in 1963, were technological spin-offs of a major development program that IBM had started in 1955. The Stretch program was intended to develop the technology to build the most advanced computers for the cryptographic and nuclear weapons communities. Los Alamos's Stretch was delivered in 1961, and Harvest, the extended cryptologic version, was delivered to the National Security Agency in 1962. Only nine Stretch machines were built, but the technology also provided the basis for the new IBM 360 computer line, introduced with the delivery of the Model 50 in 1965.

In 1957, William Norris and Seymour Cray started a new firm, Control Data Corp. The company's early products were quite modest machines, but starting with the delivery of the CDC 6600 in 1964, it was Seymour Cray and companies spun out of CDC that now defined the leading edge in high-end computing.

Digital Equipment Corp was founded in 1957 by Ken Olsen to supply transistor-based modules to the engineering and scientific communities. The company produced a computer built with these modules to process data—the programmed data processor. The PDP1, built in 1960, was the first of the minicomputers. It was an all-transistor machine with a 5 μ s cycle time, 16 384 18-bit words of core memory and no secondary memory.

At the Cambridge Electron Accelerator, a PDP1 was used to collect data from up to three experiments simultaneously. An operating system, the Time Sharing Administrative Routine, had been written for that purpose. Each experiment was assigned a 200-bit-per-inch $\frac{1}{2}$ -inch tape drive and 4096 words of memory. It took only 3 cycles (15 μ s) for the user to gain control of the processor on an interrupt, and data were transferred into memory on one-cycle intervals (1 word every 5 μ s). This was the start of a new approach to data acquisition and control, which was made especially attractive with the introduction in 1965 of the first affordable data acquisition computer, the PDP8, with a price of about \$30 000. (For the latest on data acquisition, see Joel Butler and David Quarrie's article on page 50.)

Meanwhile, with the end of World War II, many of the Manhattan Project participants had returned to academic research. Experimental particle physics research started with a return to cosmic-ray studies and moved in part to accelerators such as the Cosmotron at Brookhaven National Laboratory and the Bevatron at Lawrence Radiation Laboratory (now Lawrence Berkeley National Laboratory), as those advanced facilities came into being. Much of the earliest experimental work used visual cloud chambers or electronic hodoscopes, which are arrays that

track charged particles. The focus shifted to the bubble chamber as its technology matured after its invention by Donald A. Glaser in 1952. In these cases, whether visual or electronic, the data were mostly stored on photographic film, with one exposure per event trigger. The analysis took place off-line after the film was developed. With the advent of programmable large computers, it became possible to process a much larger number of events. The Berkeley bubble chamber group, headed by Luis Alvarez, led the way in creating mechanisms for extracting and analyzing the data from photographs. (See Robert Seidel's article on page 33.)

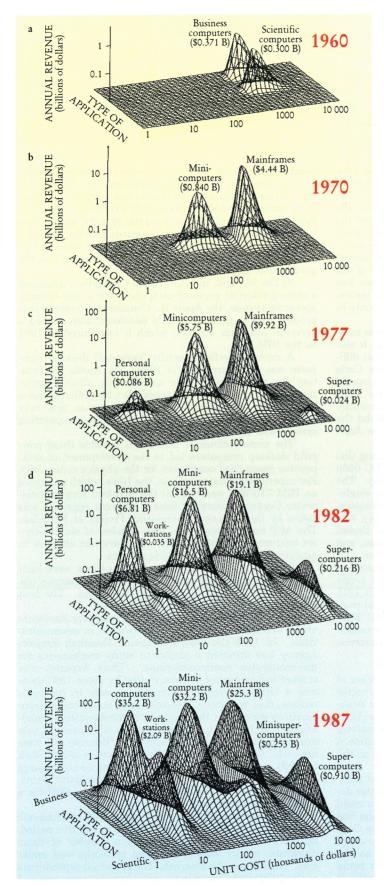
With the advent of affordable data acquisition computers in the mid-1960s, it became possible to accumulate large amounts of data electronically in real time, and partially analyze them on-line. This development opened a multitude of avenues for improved instrumentation for physics experiments, starting with the replacement of the bubble chamber with the digitally read out spark chamber.

During the 1950s and into the 1960s, much of the research on how to do computing was performed by the physics and mathematics communities, the primary computer users at the time. Some of these pioneers became more interested in studying the science and engineering of computers themselves, thereby initiating the field now called computer science.

In 1959, Jack Kilby, an engineer at Texas Instruments, constructed the first complete circuit on a single substrate of germanium. Shortly thereafter, Robert Noyce, a physicist at Fairchild Camera, introduced the idea of producing circuits using photolithography with the silicon planar technology then under development at Fairchild. This method became the standard approach for manufacturing integrated circuits and set in motion the phenomenal growth of the microelectronics industry. Year after year for the past 35 years, this industry has produced higher performing, higher density circuits at lower and lower cost, thereby making possible the information revolution.

Stable and controlled growth (1966–75)

In the ten years following the introduction of the IBM 360 series, IBM consolidated its position as the leading supplier of mainstream computers for both business and science. It was the period in which the centralized mainframe became the established mechanism of satisfying computing needs at a university, research organization or commercial enterprise. IBM supplied a wide range of compatible products and set a consistent price over the whole line. Except at the high end, no competitor was strong enough to breach IBM's very strong position. Many quite large companies, including RCA, General Electric and Philco, tried and failed—at enormous cost.



EVOLVING COMPUTER SPECTRUM, 1960-87. The annual worldwide revenues for the various computer classes are plotted as a function of the unit cost of the system and the type of application. All dollar values are given in dollars of the time. (Adapted from ref. 9.) FIGURE 2

At the specialized low end, Digital Equipment Corp controlled the marketplace with its PDP line, which was finding wide acceptance in physics and other research communities. In 1970, DEC introduced the remarkably effective PDP11. In addition to having a very elegant architecture, with a rich instruction set and a very flexible input/output bus structure (the Unibus), the PDP11 offered a choice of a number of very sophisticated operating systems. DEC's position was strong, but unlike the situation with IBM in the mainframe market, DEC had ample high-quality competition in the data acquisition and control computer domain.

Initially unnoticed by the business community was the introduction in 1964 of DEC's first general-purpose computer, the PDP6. This machine and its 1967 successor, the PDP10, were widely used as general-purpose departmental computers by many university physics departments. This minicomputer enabled DEC to compete with IBM at the lower end of the mainframe business.

By the early 1970s, data acquisition computers were an integral part of all experiments. Experiments had become quite large and complex, having been enabled to do so by the computers' ability to acquire and process large amounts of data. By the mid-1970s, trigger rates ranged from 1 to 10^3 triggers per second, and the data accumulated ranged from 10^3 to 10^5 bits per trigger, for an aggregate rate of about 10^6 bits per second. The limitation then, as today, typically was the rate at which the data could be written to the storage medium, usually magnetic tape.

Similarly, with on-line computers integrated into accelerator control systems, more elaborate diagnostic and operational tools became available, making it possible to operate more complex systems and to achieve better and more stable beams. It was possible to give each individual experiment control of the optical elements of extracted beams.

In particle physics, as in other fields of physics and other disciplines, it was the experimental community that was engaged early both in the use of on-line data acquisition computers and in the analysis of the data using large-scale general-purpose computers. The theoretical community became more widely interested in using the high-end computers only after the available computational power had increased

Computer Science versus Computational Science

t is essential to distinguish between "computer science" and "computational science." Computer science is the study of computers and computation. Computational science is that aspect of any science that advances knowledge in the science through the computational analysis of models. Like theory and experiment, it is one of the three legs of that science.

substantially. Figure 1 includes information on the computing power required to successfully tackle various problems

Initially and for some time thereafter, it was the physics and other science communities that invented the tools and approaches of data analysis and computational science. It was quite an *ad hoc* process, with little attention being paid to formal methods or a fundamental understanding of theoretical issues. Indeed, most of the developing computer science community at that time was not seriously interested in computational science issues. (See the box above.) This interest has developed only in more recent years.

But lack of interest in computational science was to be found not only in the computer science community. It also occurred within the sciences themselves. I had great difficulty in publishing a paper⁷ in 1965 on the Monte Carlo generation of elementary particle events. No physics journal considered computation to be in its province. After much discussion, the *Review of Scientific Instruments* reluctantly agreed to publish the paper. Not until mid-1966 did the first journal dedicated to computational physics, the *Journal of Computational Physics*, come into being.

Other important milestones that occurred during this period included the replacement in 1969 of the CDC 6600 by the CDC 7600 as the supercomputer of the day. The first modestly parallel processor, the Illiac IV, a single-instruction, multiple data computer with 64 processors, and the first machine to use semiconductor memory, was built by Burroughs Corp. It was designed by Daniel Slotnick, an engineer at the University of Illinois, and delivered in 1971 to NASA's Ames Research Center.

At the low end, in 1970 the first microprocessor consisting of a complete 4-bit processing unit on a single packaged silicon chip appeared—the Intel 4004. In 1975 the first attempt at a microprocessor-based personal computer, the MITS Altair 8800, based on the Intel 8080 microprocessor, became available in kit form. It presaged a major change in the availability of computing resources.

From supercomputers to PCs (1976–85)

This was a particularly rich period in diversification of available computer hardware. Introduced were the vector supercomputer, the personal computer, the scientific workstation, the "perfect" departmental computer (the VAX) and, as the period ended, the first commercial massively parallel computers. Along with these products came a maturation in digital communications. The effects on the physics community were profound.

The Cray 1, introduced in 1976, was the first commercially successful high-performance vector computer designed for computational problems in which the data could naturally be accessed as a vector. It and the evolving line of Cray Research Inc vector machines, including the Cray X-MP and the competing Cyber 205 from CDC, made possible the emergence of computational science as the solid third leg of scientific inquiry, the other two being experimental and theoretical science. Consequently, it

became possible to solve very complex problems by numerical approximation or by simulation. The theoretical science community became involved but quickly found that such computer resources were very expensive. Furthermore, the opportunity to use powerful systems was limited to those with important national security problems to solve and to the few scientists outside that community who managed to develop relationships inside it, mostly at the government laboratories.

At the low end, in 1978, Apple Computer introduced a mass-producible personal computer with a user-friendly operating system, the Apple II. Consumer interest in this product was so great that IBM decided it also needed a personal computer product, which it introduced in 1981 as the IBM PC.

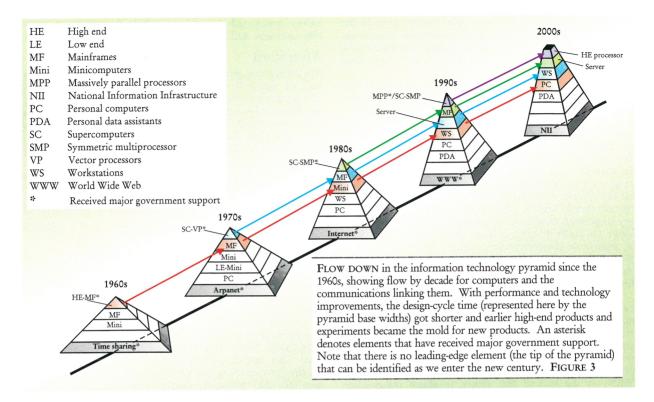
A more powerful scientific "personal" desktop computer was also introduced in 1980. The Apollo Workstation also came with a quite user-friendly, albeit proprietary, operating system. Two years later, Sun Microsystems introduced a competing scientific workstation that very quickly came to dominate the market because its operating system was UNIX, an open system.

The same technology that made possible these powerful desktop computers led to the development of inexpensive specialized computers for the physics community. For experimental data analysis, the 168/e, which emulated an IBM 370/168, was built at the Stanford Linear Accelerator Center. Somewhat later a different approach was taken by the Advanced Computing Project at Fermilab. The ACP bussed together a large number of state-of-theart microprocessor boards that were very cost effective in analyzing event data in particle physics.

Similarly, a number of parallel machines were built to perform computations in quantum chromodynamics, one of the most demanding computational areas. The best supported of these is the IBM GF11.

The rapid pace of change in the computer hardware arena was accompanied by major changes in communica-In 1968 the DOD's Advanced Research Projects Agency had initiated a project to study mechanisms for Thus Arpanet was intercomputer communications. started and continued to grow, so that by about 1980 there was a large university community, mostly in computer science, that was interconnected using public telecommunications services at the generally highest available bandwidth of 56 kilobaud. Using the ideas of Arpanet, a new academic network called Bitnet, the "Because it's time network," was initiated at the City University of New York in 1981. It grew very quickly, encompassing a large number of physics, university and laboratory participants. Bitnet was further enhanced by support from IBM for a transatlantic connection.

Similar networks were being developed by academic communities all over the world. By the end of this period, an *ad hoc* loose interconnection of a number of these disparate, mostly weakly managed networks was under way. Thus began the Internet. The parochial experiments of the 1970s were becoming a much broader cross-cultural



experiment of the 1980s.

Starting much earlier within the physics community, especially at the large government-supported laboratories, internal, very-high-bandwidth computer linkages had been developed to satisfy local computing or experimental support needs. In the mid-1960s, an evolving very-high-bandwidth intercomputer network called Octopus existed at Lawrence Livermore National Laboratory; similar systems were in place at Los Alamos National Laboratory. At all the accelerator laboratories there were broadband systems in use for accelerator control and information distribution. And in 1975 the National Magnetic Fusion Energy Computer Center was established at Livermore to support the fusion energy research community with high-performance computing. Satellite connections, initially at 56 kbaud, were established between several centers across the country.

In 1982, Peter Lax of New York University was asked to chair a panel on large-scale computing in science and engineering, sponsored by the Department of Defense and the National Science Foundation in cooperation with the Department of Energy and NASA. The primary conclusion of the panel was that although large-scale computing for computational science and engineering was vital to the country's economic and national security interests, there was a distinct danger that inadequate investments were being made to maintain US leadership.8 The panel concluded that much more powerful computers were required to solve problems of current urgency and that large and important segments of the research community were being denied access to supercomputing resources. Acting on the panel's report, starting in 1985, NSF established five (now four) high-performance computing centers. (See the table on page 26.) An important facet of this program was support for a broadband communications network connecting all the sites and their users, which evolved into the NSFnet. The physics community was and continues to be one of the major users of these facilities.

In 1985, two massively parallel commercial machines

became available. (For a description of parallel computing architectures, see James C. Browne's article, PHYSICS TODAY, May 1984, page 28.) The CM-1 was a single-instruction, multiple-data machine designed by computer scientist Danny Hillis at Thinking Machines Corp. The second was the iPSC, a multiple-instruction, multiple-data machine from Intel. These machines represented a new, less expensive approach to increasing performance, as the approach that used very fast vector processors was becoming too costly. These two companies were among the first to try to market high-performance computers based on large-scale parallelism. The problems of delivering high performance with these architectures have turned out to be much harder than originally imagined. No clearly successful path to parallelism has been discovered.

Computers and communications (1986–95)

With continued improvement in computer performance at lower cost, what had been a reasonably ordered growth suddenly exploded during this period. In the mid-1960s, Gordon Moore of Intel had observed that the number of components per chip in the most advanced integrated circuits was doubling every year, and that the expectation was for the trend to continue. It has, although the doubling time today is closer to two years. Similarly, the performance level of the leading-edge computer system over the past four decades has on the average been doubling every two years. Figure 1 shows the performance level of both leading-edge and functional yet affordable computer systems since 1955, along with some important milestones.

Figure 2 summarizes the growth of the computer business, both in revenue and in its customer range. At the beginning of this period, as shown in figure 2e, there were at least half a dozen identifiable substantial computer business classes, spanning almost the whole domain of scientific and business applications. By the end of the period, all of the functional space was covered with a good

deal of overlap, giving rise to a total realignment of the industry. The mainframe and minicomputer businesses were declining, being replaced by similar hardware now called servers. The options available to any user were now quite diverse.

At the highest performance end, the supercomputer industry is currently in a crisis. Only one US company continues to supply such machines—Cray Research, now a division of Silicon Graphics. The cost of pressing the limits of current technology beyond the 1 ns clock time with special designs has become too high to continue down that route to achieve the highest performance computers. The obvious route is the use of more extensive parallelism using components that are more readily mass-produced. But the efforts of the few companies competing to do this are not clearly viable. Although there have been some very important successes in using massively parallel machines for some grand challenge problems—for example, the quantum chromodynamics class of problems—after more than ten intensive years of study, there is still no clear understanding of how to use such architectures effectively as user-friendly general-purpose computers.

Along the way, two spectacular events occurred that dramatically changed the world. The various network experiments had developed many data and information protocols, such as Transmission Control Protocol/Internet Protocol (TCP/IP), file transfer protocol (ftp), e-mail, listserv and tools such as Gopher and Veronica. But none served as a convenient information-sharing protocol. At CERN, at the end of the 1980s, to share information among the sometimes hundreds of international collaborators involved in a single high-energy physics experiment, Tim Berners-Lee invented a data sharing approach for the Internet called the World Wide Web. Information in a wide range of formats could be accessed wherever it was stored in a universal hyperlinked arrangement. To really make the Web useful, a tool was needed to allow a user to browse conveniently. That came in 1992 with the development of Mosaic by Marc Andreessen and his collaborators at the National Center for Supercomputing

After years of effort, Albert Gore, then the senior senator from Tennessee, shepherded through Congress the High Performance Computing Act of 1991, authorizing a five-year program in high-performance computing and communications. For fiscal year 1994 the act was expanded to define the National Information Infrastructure. The technical innovations coupled with this legislative initiative set the stage for the worldwide experiment currently in progress in which the tools developed by the research community for the research community are now being tested by everyone, with the commercial world being by far the dominant player.

These advances raise potential collaborative efforts to a new level. They allow for the sharing of data and information among geographically dispersed colleagues on an almost instantaneous basis, with the potential of fusing data from quite disparate databases. The new technology allows for the remote use of instruments, the remote running of experiments and access to remote expensive resources such as special or high-performance computational facilities or digital libraries. The term "collaboratory" has been coined to refer to this new capability. 11

Major changes have occurred during this period in both the computing and communications sectors of the economy. Driven by the continued technological advances and the attainment of a quantitative level that allows for qualitative changes in the way these technologies affect society, the two sectors are transforming and coalescing. This development is leading to new options in the ways in which

society works and the ways in which the physics community interacts and does its experiments and computations.

Afterword

Reviewing the growth of the computer industry since its earliest days, it is quite clear that advanced product development to satisfy the research community has trickled down to become the basis for mass-produced products for a very much larger community. Figure 3 diagrams this technology flow-down. Note that in almost all cases, government funding has played a major role in conducting leading-edge experiments and in developing products. And the leading edge has been the engine that has generated an industry approaching \$500 billion per year in sales, an industry critical to the country's economic and national security. (See reference 12 for a discussion of these issues.) The spectacular changes in the ways in which society functions that have been wrought by the information revolution will continue into the next century. Organizations that most avail themselves of the new work paradigms will have major advantages over those that do not. That will be true in the research environment as well as in industry, commerce, education, entertainment and so on. But there are some dangers.

Throughout the cold war, the government used various mechanisms to support the development of high-end computing. Many things changed with the end of the cold war. One consequence has been a lack of understanding on the part of Congress and the American public of the spin-offs that naturally derive from the innovative approaches the research community takes to solve its problems. Now that information technology has become a very big business, it is widely assumed that the marketplace will support all the necessary innovations for the future. This assumption does not work well at the very high end. There, the natural time constant for investment and development is too long for any company to undertake in the current business climate. If this problem is not solved, then it may very well affect the leadership that the US has maintained heretofore as we enter the information age.

References

- 1. A. Toffler, The Third Wave, Morrow, New York (1980).
- 2. A. W. Burks, A. R. Burks, Ann. Hist. Computing 3, 310 (1981).
- For a review of some of the early computer history, see R. Serrell, M. M. Astrahan, G. W. Patterson, I. B. Payne, Proc. IRE [Institute of Radio Engineers] 12, 1039 (January 1962).
- 4. For a review of the early computer history for cryptologic organizations, see S. S. Snyder, Ann. Hist. Computing 2, 60 (1980).
- 5. For the fascinating story of the development of the IBM 360, see two articles in Fortune, September 1966, p. 118, and October 1966, p. 140. For a more technical perspective, see B. O. Evans, Ann. Hist. Computing 8, 160 (1986).
- 6. A. E. Brenner, IEEE Trans. Nucl. Sci. NS-12, 241 (1965).
- C. A. Bordner Jr, A. E. Brenner, E. E. Ronat, Rev. Sci. Instrum. 37, 36 (1966).
- 8. Report of the Panel on Large Scale Computing in Science and Engineering, Department of Defense and National Science Foundation, Washington, DC (1982).
- 9. IEEE Scientific Supercomputer Subcommittee, Computer, November 1989, p. 57.
- High Performance Computing and Communications: Toward a National Information Infrastructure, 1994, Committee on Physical, Mathematical and Engineering Sciences; Federal Coordinating Council for Science, Engineering, and Technology; Office of Science and Technology Policy, Washington, DC.
- National Research Council, National Collaboratories: Applying Information Technology for Scientific Research, National Academy P., Washington, DC (1993).
- 12. National Research Council, Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure, National Academy P., Washington, DC (1995).