My Second Computer was a UNIVAC I: Some Reflections on a Career in Progress

George Sadowsky

[This article is based upon the reminiscences given by the author at a party celebrating his retirement from NYU on November 17, 2000] [1]

The first time I came to New York University, it was to program a computer. It was during the fall of 1958, and the computer, an IBM 704, was installed in what is now Meyer Hall at 4 Washington Place, behind the translucent glass blocks on the second floor. The computer belonged to the AEC (Atomic Energy Commission), now known as the Department of Energy, and was installed at NYU for use by the Courant Institute.

I had recently left the graduate mathematics program at Harvard, deciding that I didn't want to be a pure mathematician, which was the only respectable flavor they turned out at the time. After spending a summer in Cambridge, doing as little that was serious as I could, I joined the computing group of Combustion Engineering (now the ABB Group) in northern Connecticut.

The previous summer, I had learned to program an IBM 704 computer during a summer job at the General Electric Research Laboratories in Schenectady, New York.[2] Almost all programming at that time was done in assembly language, with corrections often made in machine language. Operating systems consisted of fairly rudimentary input-output packages, program loaders, and core dump programs. The next fall in graduate school, I took a course in numerical analysis, and used a UNIVAC I that Harvard had managed to obtain as a donation from the Sperry Rand Corp. The UNIVAC I was an interesting machine to program, with its mercury delay line storage and its short mean time to failure. Programs were entered into the computer by typing them onto steel magnetic tape, a major innovation at that time.

Serious computing work at Harvard in the late 1950s was done on the Mark I and Mark IV computers. Howard Aiken, who ran the Computation Lab, viewed computers as research tools that were not appropriate for mundane applications. He was using the Mark IV in an attempt to generate adequate human language translation programs. Harvard's subsequent lack of commitment to general user computing can be traced to that initial orientation. The Mark I, which had a transfer bus consisting of a thick rotating steel rod some 50 feet in length which transferred the contents of memory locations that looked like vehicle odometers, was already outdated by that time.

Working with the IBM 704 at NYU was an entirely different experience from the UNIVAC I. It was built for executing scientific applications, and had as its major innovation a magnetic core memory, replacing the Williams tube memory of the IBM 701 and the primary memory drum of the IBM 650. It also had a floating point arithmetic unit and index registers to form effective addresses, both of which were significant advances at that time. The machine had the equivalent of 128 KB primary memory, 32 KB of secondary drum memory and magnetic tapes that held 5 MB of data. It operated at 0.04 MIPS, and cost $3 million dollars in 1957, the equivalent of perhaps $20 million today.

The Atomic Energy Commission had installed the IBM 704 at NYU because of the large amount of contract work in applied mathematics that the Courant Institute was doing for the Commission. At that time, largely because of the very high cost of computational facilities, wasted computer cycles were not tolerated,

---

[1] The author is Executive Director of the Global Internet Policy Initiative, sponsored by the Center for Democracy and Technology, and is a consultant to various other organizations. The views expressed in this article are those of the author, and do not necessarily reflect the views of any of the organizations with which he is affiliated.

[2] General Electric was one of the leaders in scientific computing at that time, along with Westinghouse and most companies in the then-booming aerospace industry. The SHARE Users Group had formed only a year or two earlier, the driving force being the development of interoperable standards and shareable code. This was absolutely essential at the time, given the cost of developing error-free code. The IBM 704 was delivered with no software from IBM whatsoever, although IBM quickly recognized the need for access to it and established a library where organizations could donate programs. The library was maintained on punch cards at 590 Madison Avenue in New York City.

so that other AEC contractors were given blocks of time on the system as part of their contractual arrangement depending upon their needs. Combustion Engineering, historically a maker of boilers for steam-powered electrical generating plants, had recently entered the nuclear energy business and had already manufactured several reactors for the U.S. Navy submarine program. The company was rapidly establishing itself in both the civilian and military sectors of the nuclear industry, and had a lot of design contracts with the AEC. As a result, the entire second shift of the 704 was often used for production and debugging work. My job required me to spend every third week or so in New York City using the computer, which gave me an interesting introduction to NYU and Greenwich Village.

The computing environment at that time was both primitive and exciting. Although Fortran II had just been announced and looked interesting (the original Fortran I had no subroutine capabilities), most work was done in assembly language because the Fortran compiler was buggy and slow. In addition, the initial compilation algorithms produced code that was obtuse and slow running at best, and sometimes just plain wrong. Not only did we have to debug our own programs, but we had to help find errors in the compiler as well. Since even the assembler was costly to use, corrections were often made by adding binary correction cards directly to the end of the object program deck, and sometimes so may corrections were made—some documented, some not—that it wasn't clear what program specification was executing.

The IBM 704 was, in effect, a very expensive personal computer. Only one person and program could use it at any given time, and often program execution, as well as dumping the evidence of failure, was done at the system console as fast as one could figure out what to do. The need for tools to assist in these processes was great. I remember writing a program disassembler that would take the substantially modified object programs, together with a symbol table, and produce a reasonable assembly language listing of the modified code. For some jobs I wrote a one-pass octal assembler because the IBM assembler was too cumbersome. Finally, I wrote an operating system that would do batch initiation, automatic job sequencing, system resource management, error recovery, and accounting in order to replace the constant human intervention that was then required and that caused many errors. These projects were necessary, but they also provided a great deal of enjoyment since we were able to become involved in interesting system projects for which we had no real model.

Most people only dimly perceived the future of computing at that time. I remember having several discussions regarding whether digital or analog computing was the direction of the future. One could have expected such conversations in an engineering company, but industry leaders were just as confused about the size and direction of the market that they were creating. In 1954, IBM committed to build only 18 IBM 701 systems because, in an opinion attributed to Thomas J. Watson, Sr., that number of systems would be able to satisfy the computing requirements of the world for the foreseeable future. Watson was not alone in expressing such limited vision; the editor in charge of business books for Prentice Hall at that time is reputed to have said, "I have traveled the length and breadth of this country and talked with the best people, and I can assure you that data processing is a fad that won't last out the year."

The hardware of the IBM 704 provided some interesting challenges. Like other systems of its generation, it had only partial instruction de-coding. While there was a 12-bit instruction field in every 36-bit word, allowing for, in theory, over 4,000 instructions, only about 100 instructions were actually defined. This structure helped the designers who implemented the execution module of the system, since they could tie specific functions to specific bits, but it left many combinations of bits in the instruction field undefined. Execution of any such combination was at your own risk. However, after the machines had been designed and built, combinations were discovered that were actually very useful—the one I remember best is the STZ (Store Zero) instruction, which cleared a memory location. It was not part of the original instruction set, but was discovered later, and then, of course, quickly added to the assembler op-table and frequently used.

I remember writing a core dump load module that depended upon being able to preserve as much of the machine state as possible while loading itself into the system. It was a torturous process, but depended crucially upon a combination of bits that were undefined. Max Goldstein, who was even then in charge of the computation center, persuaded me to work with one of the two permanent on-site IBM engineers to determine what it did. So, we spent quite a few hours poring over circuit diagrams, tracing pulses and gate settings, and finally determined that this combination of bits did multiple operations that we absolutely needed to make the dump initiator work.

The environment of partial instruction decoding was not all positive. Since both program and data were in the same unprotected memory space, a program addressing an error could cause an undecipherable instruction to be executed with unknown effect, leading to some long and ugly debugging sessions. In such instances, as in others, the best way to proceed would often be to dump to the printer the entire contents of memory and then become a sleuth to unravel the mystery of the observed system behavior.[3]

Although we did not mingle much with the Courant folks in those days, we did know some of them. Occasionally, I would discuss a computing issue with Max. I remember Henry Mullish working on the system, and I can recall seeing both Peter Lax and Jack Schwartz from time to time. Florence Ragusa was, at that time, one of the more accomplished system experts dealing with the 704.

Sometime in 1960, the main offices of Combustion Engineering moved from New York City to Windsor, Connecticut, where the Nuclear Division was headquartered. The Division had previously put a big bet on the civilian reactor business, and had hired Walter Zinn from Argonne, a man who had very bad relations with Admiral Rickover of the U.S. Navy, to lead the way. It was the beginning of the long, slow demise of that line of business for Combustion, as the civilian reactor business was beset by a myriad of problems. The main offices moved into a large building vacated by the Nuclear Division, and one of the people who made the move was Edi Franceschini, who had temporarily left the life of the arts to immerse himself in the real world. "Big teakettles," as he referred to Combustion's boiler business, seemed right in the middle of the real world to him. Those of you who know Edi will appreciate why, as he more fully understood the path that he had taken, he retreated from that business into the world of computing and abstract mathematics.

After leaving Combustion Engineering in 1962, I went on to Yale University, first to work, and then to attend the graduate school. After that, I moved on to work in Washington and then to the United Nations in New York to do computer technology transfer work in developing countries (a thread that would reassert itself in interesting ways later in my career). In 1986, I left the UN to go to Northwestern University, and after four years in an environment that seemed parochial compared to New York and the rest of the world, I was looking forward to the opportunity to return to a more advanced and cosmopolitan environment. The opportunity arose in 1990, and I returned to New York University, this time to take over the Directorship of the Academic Computing Facility from Max Goldstein, who had handled most of NYU's computing needs for over 30 years.[4]

Quite understandably, much had changed at NYU since 1958. The Courant computing group had formalized itself as the Courant Mathematics and Computing Laboratory (CMCL) in 1964, as Warren Weaver Hall was being constructed, and had installed serial no. 4 of the CDC 6600 as its computing platform. By the time I arrived, Control Data computing systems had just been phased out of use, and there was a plethora of systems being used for various purposes. During my first year back, I counted 16 different flavors of UNIX being run in the Warren Weaver computer room.

The CMCL was not initially given any responsibility for serving the general computing needs of the University, so in the 1970s a separate group was established to serve those needs. I remember visiting the facility in the late 1970s and observing an overloaded IBM 370/145 trying to cope with the load, while faculty and students tried just as hard to cope with some variant of JCL. For whatever reasons, the independent academic computing service did not succeed, and Max and the CMCL were asked to take

[3] I remember the first time that I had to do this, at General Electric in the summer of 1957. The only output device was an online printer that generated 150 lines per minute. When I had dumped the entire memory, consisting of 8KW (=32KB), I calculated that the machine time used for this one operation cost more than my entire week's salary. And it didn't help that when I found the error, it turned out that I had used an integer ADD instruction instead of a floating point FAD instruction. I don't think I ever confessed to it, given the cost of finding it.

[4] Max's tenure at NYU was very fruitful. After coming to NYU from Los Alamos, where he had worked during the second world war, he established and directed Courant's computing activities for more than 30 years. He was responsible for establishing the Courant Mathematics and Computing Laboratory (CMCL) when Warren Weaver Hall was built in the early 1960s, and he, along with Jack Schwartz, were the primary founders of the Computer Science Department at NYU.

on the responsibility in 1981, resulting in the formation of the Academic Computing Facility (ACF).

In 1990, the ACF still very much resembled the earlier CMCL in orientation. Large machines dominated, while the desktop was largely ignored. The ACF had passed up the opportunity to involve itself in microcomputer resale activities, and was a bit reluctant to get into the student microcomputer laboratory business. Apparently, at NYU real programmers did not use microcomputers, a mistaken judgement that did not help the evolution of the ACF. Use of ACF facilities was still dominated by scientists and mathematicians, although several staff members, notably Ed Friedman, had made serious forays into other areas of application. Tellingly, the User Services Group was still classified as a subset of Technical Services, which gives a clue as to the mental model underlying the ACF at that time.

One very fortuitous development in the evolution of the CMCL was their early involvement in computer networking. In the early days of scientific computing, there were different types of computers, all having distinct operating systems, programming languages, and data formats. Interoperability was the exception, not the rule. The post-World War II climate was quite favorable for the growth of scientific research, and applied mathematics research groups sprung up at other research centers and universities. Much of this research required collaboration among this community and often required the use of distant computing facilities to implement such collaboration.

This requirement was at the heart of the development of the initial ARPANet. Using file transfer and remote login techniques, it would be possible for someone at one location to use a computer at another location as if they were physically present. Because of the very large research productivity gains that this mode of operation promised, mathematicians and computer specialists became very excited about achieving this goal, and the initial ARPANet results spurred them on. Courant was one of the leaders in this field, and the CMCL staff, notably Edi Franceschini and Bill Russell, contributed substantially to parts of the ARPANet protocols and applications in the 1970s and early 1980s.

In 1981, when the Courant Institute was asked to take over the academic computing support responsibility for all of New York University, networking was still in an experimental stage. The mode of development seemed to be towards disciplinary networks such as CSNET for computer sciences, HEPNET for high-energy physicists, and BITNET for mail and file transfer. This was possibly encouraged by ARPA, which, although it had the closest thing to a general-purpose network at that time, restricted its use to those working on government and military activities. However, one should note that the management of ARPANet at the time had no concept of how large their network would become, since the pre-TCP/IP protocol set, NCP, had space for addressing no more than 256 networks of computers.

In 1983, ARPANet made a very significant transition from its earlier protocol set, NCP, to TCP/IP. It was thought that the theoretical capacity of TCP/IP to address four billion hosts would suffice for the indefinite future, exemplifying the sense at that time that the Internet was likely to remain a relatively specialized network.

There were other competing networking efforts at the time. UUCP (Unix-to-Unix Copy Program) was used across dial-up links to implement a store and forward network for delivering mail files. Similarly, FidoNet was a network created by a grassroots effort both in the K-12 community and by non-governmental organizations (NGOs) operating in developing countries to develop a cooperative volunteer store and network using dial-up telephone circuits. At NYU, there were some incursions by specialized disciplinary networks, but the CMCL under Max's leadership stuck to the general-purpose network, and as a result, was on the forefront of offering the best possible network services, whatever their experimental state, to the NYU community.[5]

The first construction of networks was motivated by the incompatibility of resources and the need to use computational resources that were distant. Remote login and whatever followed the login were the driving features. However, it wasn't long before the communication features enabled by the network grew and became more dominant. Electronic mail may not

[5] Max's tenure at NYU was very fruitful. After coming to NYU from Los Alamos, where he had worked during the second world war, he established and directed Courant's computing activities for more than 30 years. He was responsible for establishing the Courant Mathematics and Computing Laboratory (CMCL) when Warren Weaver Hall was built in the early 1960s, and he, along with Jack Schwartz, were the primary founders of the Computer Science Department at NYU.

have been the ARPANet's killer application, but it certainly grew to large proportions.

In the early- to mid-1980s, the National Science Foundation decided to invest in a national backbone and to subsidize regional networks. Its primary stated motivation was to allow researchers to obtain access to the five supercomputer centers that were being established in different parts of the country without having to travel to them. Remote access to computational resources was again the primary motivation for building the network, although the forms of remote access were improving rapidly. In the early 1990s, however, the primary motivation to use the network was shifting substantially to communication, dominated by e-mail.

The introduction of the World Wide Web shifted this balance again, causing access to remote resources to become the primary use of the Internet. The difference was that this time the remote resources were generally content based, not computationally based. It certainly took a lot of computation in the background to make the content available, but the computation could be distributed and, thanks to the continuing work of Moore's law, which states that one can double the number of transistors (and therefore roughly the power) on a chip every 18 months for the same cost, it was increasingly inexpensive to provide it.

Today, we have the best of all worlds, with many content and computation servers readily available, as well as the possibility of communicating with almost anyone who uses the Internet. In the United States, the interplay between these modes is now natural and intuitive. In developing countries, however, the Internet is a new phenomenon, and people are just beginning to learn how to mix these types of uses of the Internet to more fully exploit its resources for development purposes.

It seems to me that during the past 10 years, NYU has been burdened by ambivalence regarding the path that central academic computing should take. On the one hand, a first-rate research faculty has demanded, and generally obtained, a fairly high level of central computing and networking support. On the other hand, it seems to me that on the student side, with 45,000 students, there is always the underlying possibility of being overwhelmed by demand, or of having expectations arise that could create a new and unexpected demand with significant associated costs. I remember that in 1990, students were not allowed to use the microcomputer labs unless their instructor had requested and received a course account; as a result the labs, which had almost the same number of computers then as there are today, were largely empty. Further, centrally funded academic computing and networking at NYU has grown at a real rate of 3-5%, while the rate of investment and growth in the private sector has been several times that. I believe that we have paid for this in terms of falling behind our peer institutions in significant ways.

The computer and networking revolution has been an extraordinary and exciting one and I feel very privileged to have been a part of it thus far. Those of us who work in information technology are the beneficiaries of technological progress unmatched in nearly any other industry. Since 1955, the performance-to-price index for information technology hardware, a scale that reflects technological progress, has been in the range of 25-30% per year. Historically, the same measure for the telecommunications industry has been considerably lower. In recent years, however, as this technology has become increasingly dependent on the semiconductor industry and digital communication using optical fiber media, technological progress has accelerated sharply.[6] Moore's law clearly still holds. Dense wave division multiplexing techniques and optical switching technology promise substantial short- to medium-term decreases in the cost of wide area telecommunications. And there is no end in sight, at least not for the next 10 years, which is generally the limit of product development and engineering vision.

Though the cumulative effect of this rate of progress can be measured quantitatively, it is felt qualitatively, as applications that were once infeasible suddenly become possible. Markets spring up without much warning, the demographics of capitalism continue to accelerate, and firms are born and die off with great rapidity. We have appropriately chosen to call this phenomenon Internet time, and it sets the speed of the clock for those of us in this profession.

The nature of academic computing has changed substantially during my association with it. In the early 1960s at Yale University, when I managed the Yale

[6] Recently there was an announcement that SURFNet in the Netherlands had been connected to the STARTAP in Chicago with 32 lambdas (light wavelengths), each of which has a raw capacity of 10 GB/second. Such capacity would have been unimaginable 10 years ago, but is likely to seem commonplace 10 years from now.

Computer Center, the potential of the applications space for digital computing was relatively unknown. This was an exciting time for quantitative techniques in general, since algorithms that had little hope of being executed manually could now be implemented for large scale execution on digital computers. Operations research techniques took on new meaning, and there was extensive development of linear programming and related techniques. I believe that this was, in part, responsible for the Soviet Union's faith in the ability to centrally plan its economic progress using such techniques, a faith that was unjustified in hindsight, and not because of any lack of computational power.

At that time, the universities were one of the few sectors empowered with high-speed scientific computing, and they were a magnet for businesses and entrepreneurs wanting to discuss and implement experimental techniques. Programs were often proofs of concept rather than implementations of specific requirements analyses. There were, of course, commercial computing operations in most large companies at the time, but they were more concerned with the automation of large file organization and processing. Many useful innovations flowed from these activities, but commercial computing was viewed as a different culture that chose to work on a different set of interesting questions. At that time, there were two IBM user organizations, SHARE for scientific computing, as it was known then, and GUIDE for commercial computing. In an attempt to identify the strengths of commonality among members of the groups and their organizations, a joint meeting was held in Atlantic City in the late 1960s. A vote to merge the organizations that was put to both memberships shortly after that was voted down by a wide margin. The cultures were not compatible.

With the passage of time, academic computing has slowly but inexorably shifted toward more of a utility or service model. The ability to exploit computers on a simple level has spread from a relatively small group of technical people to a substantial proportion of the world's population. University students, staff, and faculty now have the benefit of 40 years of experience with which to assess the extent to which, and how, computers can help them in their work. For many daily activities, the loss of computers would result in catastrophic losses of productivity. By and large, innovation in information technology has been absorbed within universities by the faculty and staff, and outside the universities by the industrial sector, particularly the startup sector of the industry. This would not have been possible without the continued functioning of Moore's law, which lowered the barriers to entry into the industry by orders of magnitude over a relatively short period of time.

Thus, academic computing as a named field is now disappearing. The name "Information Technology Services," now used at NYU, better captures the major role of information technology not only at this university but also at other institutions. Our old name, the Academic Computing Facility, became a complete misnomer during the 1990s. As we grew in the areas of both networking and *e*Services, our role transcended the purely academic. Computing became only one of many services we offered. And finally, although facilities were still an important part of what we provided, we could have provided similar services without running an actual facility.

The current paradigm, providing services through the creative and productive use of various information technologies, suits the University's needs. This does not mean that innovation has become unimportant; it does mean, however, that the locus of that innovation in the academic sphere has shifted away from the previous technology providers toward faculty members and their collaborators.

In developing countries, however, one obtains a very different view of the position of computing and networking, certainly in universities, but not only there. In April 2000, I was working in Madagascar for several weeks, helping USAID prepare a plan to connect five major universities in that country to the Internet. The Leland Project, a USAID project to bring Internet connectivity to twenty African countries, had succeeded in providing a moderate bandwidth Internet link to Madagascar in 1997 via VSAT, and already there were eight Internet service providers (ISPs) with several thousand customers.

The universities, however, did not have the substantial funds that were required to afford even a 64 Kbps link to the Internet, which would have cost each of them about $20,000 per year. Most of them had a single dialup connection at 28.8 Kbps to a provider. The universities ranged in size from less than 1,000 students to 15,000 students. Many members of the faculties had been educated in European Universities in places such as Paris, Grenoble, Lyon, and Lausanne, and understood exactly what it meant to have access to information, and were deeply frustrated by the lack of access. Some students also understood what the Internet could provide. It was a stark contrast for me to hear the passionate arguments in Madagascar for access to knowledge via the Internet, having recently heard the passionate arguments of NYU students for access to Napster via the Internet.

There is a very large world outside of NYU, and given the responsibilities and amount of activity that we have here, it's sometimes easy to ignore or forget that and to become parochial, even in the midst of this rich cosmopolitan environment. I was fortunate enough to be able to work for 13 years at the United Nations in the 1970s and 1980s and participate in the transfer of information technology to developing countries all over the world, and I've continued to do that during my years at NYU. I've worked in about 40 countries, almost half of which are in Africa. I've also spent nearly 6 months in the People's Republic of China helping to process the 1982 Census of Population and Housing. It's a very rich and diverse world, and being able to experience so much of it has allowed me to put my cultural outlook and that of the United States into a much richer global perspective.

Shortly after arriving at NYU, I joined with other volunteers and the Internet Society to initiate a program of network training for people from developing countries. The training program encompassed, *inter alia*, basic connectivity training, advanced connectivity training and routing protocols, discovering knowledge and providing information services on the Internet, national network management, and building and managing viable ISPs in developing environments. We've held workshops in California, Prague, Honolulu, Montreal, Kuala Lumpur, Geneva, Yokohama, Rio de Janeiro, Mexico City, Mérida, Bamako, Budapest, Warsaw, Yaroslavl, Ohrid, and a number of other cities in eastern Europe to train about 2,500 students in these topics. By doing so, I think we've accelerated the introduction and exploitation of the Internet significantly in those places where it was most needed.

The issue of inequality among people with respect to Internet access is now achieving substantially increased recognition. It is often referred to as the digital divide. The World Bank has a number of programs targeted in this area, and the G-8 nations focused on this issue at their meeting in Okinawa in July 2000, leading to the creation of the Dot Force initiative. The UN Secretary-General has set up a task force to study this problem led by Jose Maria Figueres, the former president of Costa Rica. Even the MIT Media lab and the Harvard Center for International Development have joined forces to establish the Digital Nations Consortium to address these problems from a primarily technological viewpoint.

The phrase 'digital divide' is, at best, a confusing code word for the disparity between people who have enough computing and networking resources and those who cannot obtain sufficient access to them. It is an expression born out of the problems of poverty, both in this country and in the developing world. The conventional wisdom is that there's a great gap between those who have such resources and those who don't, and that it should be closed as quickly as possible. The expression is often used in such a way as to imply guilt and embarrassment on the part of rich countries, so that in addition to being a technical challenge, its elimination becomes a moral imperative also.

I believe that it's a mistake to regard the digital divide in this manner, since it reverses symptoms and causes. At best, the digital divide provides some evidence regarding our success in dealing with other divides. At worst, it's a phrase that should be banned from the language. The way in which it is often used, i.e., closing the digital divide, suggests that the important issue here is to provide computers and network connections, that those actions will solve the problem. In my view this solves the wrong problem, if it was in fact ever a problem to be solved.[7] The real divides that afflict the world are more basic. We have economic divides, educational divides, social divides, divides in our ability to obtain adequate health services, and divides in ability to find and keep jobs. These divides are exacerbated by racial divides and divides regarding sexual preferences that confound the situation further. Such divides exist between countries, within countries, regions, and most communities. Furthermore, they aren't so much divides as they are distributions along a spectrum, ranging from very satisfactory to very unsatisfactory.

[7] In a strange twist of logic, Seymour Papert suggested at the Digital Nations Conference at MIT in October 2000 that much of the problem of the digital divide would be successfully addressed if schools in developing countries were given computers, even in the absence of any network connections!

Rather than looking at digital divides as a problem to be addressed, we need to look at the real divides, and how information technology can assist in narrowing those divides. The conventional wisdom is beginning to shift toward talking about digital opportunities, indicating a change in focus. Information technology is a powerful tool, and if we don't succumb to the hype, and instead focus upon real costs and benefits, it seems clear to me that substantial investments made in the general context of addressing the real problems can yield major returns. If we can address these concerns successfully, the symptom of the digital divide is likely to disappear.

The recent rise of distance education is a phenomenon that intrigues me. Distance learning has, of course, a long history, going back to correspondence courses that were common prior to the invention of the computer. However, the distance education efforts of today are more properly the latest stage in the evolution of what started as computer-assisted instruction, or CAI.

Education is, of course, a broad topic. In its broadest sense, it encompasses learning both inside and outside of any formalized instructional structure, and even encompasses the experiential learning that we can perceive every day. It's also a very important topic; more than 30% of people in the United States are fundamentally involved with education, either as teachers, students, or staff. In the western world, we recognize the connection between education, research and economic progress.

Computer-based education has existed for as long as computers have. Educational technologists have been keen to exploit any real educational opportunities that technology provided. Of course, not all such opportunities have lived up to their initial promise. Television is one good example of this. Once thought of as a near panacea for education, it has evolved into a large entertainment and information medium, with educational television occupying a niche position.

Early large batch processing computers were used for science, engineering and management instruction, among other things. The first computer specifically directed to educational instruction was the IBM 1500, a variant of the IBM 1130 that was introduced in the mid-1960s. The system had a course development language and implemented programmed instruction logic, but was a commercial failure. The largest project I'm aware of in this field was the Plato project at the University of Illinois in the 1960s and 1970s, headed by Donald Bitzer. Bitzer developed a special purpose machine and the plasma display tube to implement Plato courseware. It was a strong and massive effort, but the costs of both the station and course development were too high to be commercially feasible.

A version of Plato was implemented by the University of Illinois Plato team to the IBM PC/XT and PC/AT in the early 1980s, but was not successful in the marketplace. Later, when microcomputers began to penetrate university environments, both IBM and Apple set up repositories for sharing educational software. Although some programs were of moderately widespread use, the experiments were generally failures. The advent of inexpensive CD-ROMs for recording and distributing content has been more successful.

Each of these computer-assisted education technologies has yielded benefits and has met some set of needs. However, they have generally satisfied specific niche markets as opposed to being general purpose, with the possible exception of CD-ROM content distribution. There have been some real successes, e.g., Mathematica for mathematics teaching, and drill programs for language instruction. In addition, general productivity applications such as word processing, spreadsheets, presentation and database packages have clearly had an intermediate effect upon the productivity and quality of the academic process, but it has been difficult to measure.

There are a variety of reasons why these efforts were not very successful. One of the most important was the accumulated experience indicating that it took about 100 hours of courseware preparation time to prepare one hour of classroom or study time of significant quality. This ratio of 100:1 was forbidding for all but the hardest working faculty members, and the market sizes for the products were not sufficiently large or profitable to entice commercial participation. There was also the issue of platform. With multiple microcomputer platforms and configurations and multiple operating systems whose lifetimes were not predictable, the market for any one product was fractionated. Finally, issues of updates, and physical inventorying and shipment of products added to the complexity of providing acceptable service. Other issues, such as copyright and determination of ownership, further clouded the prospect of decent returns on investment.

Given this history, what can we say about the prospects for success of distance education, perhaps equally aptly named network-assisted learning? It's a very important question because, if successful, it promises to do a great deal to narrow some of the divides mentioned above, both in this country and elsewhere in the world. Many of the difficulties that plagued earlier efforts have been substantially reduced or eliminated. We now have a common platform in the web browser, with the distributed World Wide Web in the aggregate as the content server. There is sufficient standardization in the underlying HTML and derived standards to ensure that there is just one market. Client system platform heterogeneity is not an issue. Nor are previous problems of physical inventories, version updates or bug fixes, since they all occur on the server side and are easily controllable.[8]

In addition, during the last few years comprehensive authoring systems have been developed which allow much of the labor of writing courseware to be minimized. These authoring systems are in widespread use and appear to be commercially viable. Also, for each course developed, not only is that course's specific material available, but there are also an enormous variety of content resources on the net that can be linked in at no charge. Further, since the Internet is sufficiently pervasive and effective for a large class of learners, transport and delivery of material has almost disappeared as an issue. Finally, this reach of the Internet means that every product is capable of competing in and attracting a worldwide market of student and institutional buyers, since distance has largely been removed as an impediment to the delivery process. Altogether these represent a very major improvement for this industry.[9]

Universities have welcomed distance education possibilities with open arms. Many are entering this market, either with their own course offerings or with the help of for-profit affiliated companies or consortia. The motive seems to be some combination of increasing their catchment area, achieving increased market share, making more money, and reaching students who could not otherwise study at the university. Based upon press reports, there is a great deal of investment in such programs, and it is possible that the aggregate investment dwarfs the current demand for education that can be delivered in this manner. There is the possibility of a soft landing for some of the industry, not unlike the dot.com fallout that is happening at present. Several distance education programs have gotten off to a poor start, and one system, the California State University system, discontinued its efforts a while ago.

I think that this is natural for any new industry. The market for distance education, or network-assisted learning, is new and relatively unexplored. It seems to me that past experience will dominate in the long run. We will discover that this technology is a niche technology, and that it satisfies some needs very well and others poorly. Our hope is that the niches will be large and important, and that we will be able to use the technology in an increasingly effective way in those niches to deliver a large amount of quality education to a lot of people. The results clearly will not be available for some time.

Here are some tentative predictions regarding the future of network-assisted learning (NAL):

- • NAL will make the delivery of courses within universities considerably more productive, as faculty members move to take advantage of online syllabi, notes, reading materials and other course information. Course material will become reasonably rich over time, enhancing the learning process.
- NAL will empower the active learner and will reward the curious. The learning process will

[8] The rapidity with which content can be updated does have adverse effects of a new kind. In a world where content is not necessarily stable, references to content can become inappropriate, incorrect, or broken. This is of serious concern to the library community as well as others. I am reminded of various scenes in George Orwell's novel, 1984, in which history is blatantly rewritten as fast-moving political alliances shifted and reversed. While Orwell's use of the technique was a satire on the Soviet penchant in the late 1940s to furiously rewrite history according to communist doctrine, it is possible that the current potential fluidity of the information stock could be used for similar purposes.

[9] Editing note: This manuscript was written and published just before the path-breaking announcement of the Massachusetts Institute of Technology that it would migrate its entire curriculum to the web over the next 10 years. At its best, this means that any student in the world who had access to the web and who understood English would have access to the same curriculum as a student at M.I.T. the implications for teachers in developing countries is perhaps even more dramatic, for they obtain a curriculum that they know is correct and that they can then use to teach from locally. It remains to be seen what effect M.I.T.'s announcement will have on the distance education market place. I believe that it will be quite significant.

benefit more from this innovation than the teaching process. For developing countries this is crucial, because the students who can make best use of these resources will be able to self-identify much more easily and will be able to exploit the new resources without necessarily being first identified by the educational system.

- The balance of success using NAL will tilt toward subjects that could more properly be classified as training, i.e., education that is focused on learning specific skills. Classical university education will also benefit from NAL, but in a manner that supplements traditional educational models rather than substituting for them.
- NAL will both facilitate and increase the requirements for continuing professional education, and to a lesser extent testing, among well-educated professional groups in most countries.
- NAL will succeed in the developing countries and regions of the world only to the extent that the national gateways and the network backbone, regional and tail circuit infrastructure will support the aggregate level of traffic needed to make it viable.[10] NAL is one of the most important reasons why this infrastructure needs rapid and radical improvement in most such countries and regions.

The niches in which NAL best meets demand will be determined by the market. This is perhaps the first technology that is likely to be cumulative, so that both curricular and other material and experience will aggregate over time. As a result, existing niches are likely to grow and new niches to open up as technology changes.

Effective global governance and management of the Internet is critical to network-assisted learning, to the continued migration of useful content to the Web, and to the success of just about every other aspect of the emerging information society. We all need to concern ourselves with it. I don't believe that there is any other phenomenon like the Internet that has grown so quickly and is essentially self-regulating. Its technical standards are defined by the work of the Internet Engineering Task Force (IETF), an open membership organization that operates as a meritocracy and is ICANN's Protocol Standards Organization. Internet addresses are allocated by a group of regional registries, which in turn are governed by their members, which include the major Internet Service Providers in their region.

It is essential to make a distinction here between Internet governance and Internet management or administration. Governance is going to be supplied by governments, or by international organizations as their proxies, no matter what others think. They will apply their laws and exercise their role vis-à-vis the Internet. Governments will concern themselves with possible illegal activity—civil or criminal—when any protected right to confidentiality may be violated, with the commercial regulation and taxation aspects of Internet commerce, and with any issues regarding national security. In addition, governments will also concern themselves with intellectual property and trademark issues, but jurisdictional issues may well arise prompting use of an international venue for dispute resolution.

Governments will also concern themselves with cultural threats posed by the Internet, although in my opinion the combination of satellite television and the Hollywood film industry pose a graver threat to cultures having values significantly different than American values. However, each country has its own culture and the behavioral standards that derive from it, such as its view of the right to "free expression," its tolerance or intolerance of specific content such as sexual content and pornography, and its position on religious liberty, or the absence thereof. All of these areas are the legitimate concern of the State, and will be exercised by the State; national positions regarding these issues are woven into the code of justice and the culture and customs of the country.

International organizations, which obtain their legitimacy through national governments that act as their sponsors, are also interested in the Internet. The International Telecommunications Union (ITU) is increasingly eager to play a part in the governance of the Internet. The Organization of Economic Cooperation and Development (OECD) is interested in the conduct of international commerce on the Net. The

[10] Designers of instructional materials should share this responsibility to some extent, recognizing that much of the world is connected with "thin pipes," and will not be able to use content that required exte4nsive bandwidth. Unfortunately, designers are generally responding to market pressures, and attractive delivery is often equated with extensive use of bandwidth-consuming multimedia. This eliminated large markets of admittedly remote and generally poor people who could benefit from receiving it.

European Union has a special interest in the confidentiality of information about individuals on the Net. Finally, the World Intellectual Property Organization (WIPO) is working to resolve conflicts regarding domain names and trademarks on an international basis.

So, what functions remain for the self-regulation of the management of the Internet? There is general agreement, to which I subscribe fully, that international organizations in their present form and functioning could not manage the Internet even if they wanted to (which some of them want to do). The success of the Internet comes from its structure and evolution, which has relied upon bottom-up development and decentralized management. International organizations are not capable of implementing or managing such a process. We need an organizational structure, most likely not-for-profit, that will allow us to do this. The goals that we want to achieve are the following:

- to share the IP address space over regions and countries in an equitable manner
- to evolve the available resources and generate additional resources as required, such as the current evolution from IP version 4 to IP version 6
- to define the standards that continue to ensure interoperability among the many networks that make up the Internet
- to manage a process that defines the taxonomy of top level domains in a manner that most of the world agrees is appropriate and equitable

Those are our primary responsibilities, not those of the State.

In 1998, the Department of Commerce of the United States chartered the Internet Corporation for Assigned Names and Numbers (ICANN) to fulfill these functions on its behalf. ICANN took over a set of functions, somewhat expanded, from the Internet Assigned Numbers Authority (IANA), which was located at the University of Southern California and had exercised similar authority until 1998. ICANN now administers these functions through three standards groups, the Address Supporting Organization (ASO), the Protocol Supporting Organization (PSO) which is the IETF, and the Domain Name Supporting Organization (DNSO). Only the last of these has been truly controversial.

In 1998, under pressure from other groups that wanted to replace IANA, the U.S. government required that ICANN develop an organization of individual members from whom some of ICANN's Board of Directors would be chosen, without defining any sense of what individual membership should consist of. In my opinion, this was a mistake, but it has led to the exploration of some very interesting questions. ICANN interpreted the requirement in such a way as to let anyone in the world register as an ICANN member free of charge. It then held an election for five new members of the Board of Trustees, elected regionally, by the individual members of ICANN. Given that ICANN now has delegated responsibility to manage the Internet's core functions, this raises some interesting questions:

- Who really should be members of ICANN—Everyone in the world? Organizations? Companies? Individuals connected to the Internet? All of them? And from which countries?
- Who decides who represents the Internet community? If the representatives are elected, who has the right to vote?
- Under what circumstances is this even a useful concept?

In the year since ICANN decided to form a general membership, they have registered 150,000 members, of which about 75,000 registered in time to vote. Of these, at least 15,000 were Japanese, and 15,000 were Korean, the apparent result of strong campaigns in these countries to register many members. One might speculate upon how ICANN membership would develop if the Chinese government decided that China should be well represented among ICANN members. Even if the voting is by region, this behavior raises the specter that national bloc voting in the future may eventually determine ICANN Board membership. To make matters more complex, a group of dissidents largely from the not-for-profit sector, apparently not wanting to stop with just the vote that had taken place, started an ICANN individual members organization to act as a counterbalance to ICANN.

This raises the issue of how to implement any effective form of democracy on the Internet. What are the strengths and weaknesses of each form of Internet democracy? I am concerned that there is something inherent in our interactions in cyberspace that brings out the argumentative and destructive in us as well as something that brings out a desire to be productive and positive in our interactions. We seem to lack the etiquette—the interpersonal protocol, if you like—that encourages productive and cooperative behavior on the Net. Perhaps an example will make the point more clearly. When the United States was just beginning to be settled, New Englanders practiced a form of local government called the town meeting. Periodically, everyone in the town would gather to discuss subjects and issues of local interest and concern. In general, the discussion was likely to be frank, because the issues discussed had to do with the town's wellbeing, and perhaps occasionally with its survival. But the participants must also have been at least minimally polite, for it would be necessary for them to cooperate on a day to day basis after the town meeting. In addition, they all knew each other; no one was anonymous. On the Internet, however, we observe a different mode of behavior. Cyberspace is a very large space; millions of people inhabit it at times. In addition, many of its inhabitants are anonymous; we know them only through their e-mail addresses and what they choose to reveal about themselves, which may or may not be true. Even if they are known, they generally exist at such a distance from most of the people engaging them in discussion that physical meetings are highly unlikely to occur, and can be avoided. The net result is that there is no social pressure to cooperate, as there was in the case of the New England town meeting.

Often the result of such a situation is that there is a polarization of opinions rather than movement toward convergence, and a tone of behavior that becomes less and less courteous as the process of polarization occurs. We call this behavior "flaming," and it is common in discussion groups on the Net. Those of us who follow any reasonable number of such groups are not surprised to observe this behavior relatively frequently.

In the real world, there are ways of sharing one's opinion with the rest of the world. For example, in Hyde Park in London, there's a tradition that whoever wants to speak can bring a soapbox to the park, stand on it, and give his or her opinion on any subject whatsoever for an indefinite period of time. In Hyde Park, as on the Internet, anyone can give opinions with the same degree of emphasis. However, on the Internet, if one has the time, if one can type well and quickly, if one can send lots of e-mail, and if one is articulate, then it is possible that one can convince lots of people that one's opinions are right and that a lot of people agree, even if it isn't true.

In Hyde Park, it's possible to observe how many persons are listening to any of the speakers, and to what extent there is even discussion or agreement with the speaker. On the Internet, though, it's not possible to know how many people are listening or how many agree or disagree. There simply isn't any way to do it, so generally the loudest voices determine the sense of the discussion. The anonymity of cyberspace and of the communication medium distorts our ability to measure what is really happening.

We need to find a modality of communications that works in cyberspace, and that works well enough in spite of the difficulties that this medium causes in our communications. We need this to make the management of the Internet work effectively, so that nations and the international organizations that represent them do not get the idea that they will be needed to manage the Internet. We have not yet reached this goal. As one of my colleagues remarked to me, "Parliament will never have its meetings in Hyde Park."

Finally, looking back on my career at NYU, I think we've accomplished a lot in the last 10 years. During that time, the Academic Computing Facility evolved from a scientifically-oriented computing organization to a much more general purpose academic computing support organization. Among the turning points were the creation of the Arts Computing and Humanities Computing groups, the emancipation of user services from technical services and the growth of a multifaceted help center. The separation of the information function from the technical function, and its growth, aggressive leadership and coordination in the evolution of NYU Web, were key to restructuring for the information age. The colloquium series on the use of computers in instruction and research brought speakers from other institutions that were involved in the use of computing in their disciplines and countries. When NYU was ready to move to a more comprehensive and current framework for offering information technology services, the branches of the old ACF offered appropriate support structures for significant parts of the new organization.

During my time at NYU, I am very gratified to have worked so closely with an intelligent, thoughtful, provocative group of people—the managers of the ACF and key staff people who served under them, as well as some new colleagues who joined NYU to help form Information Technology Services. By and large, we worked together as a team, often successful, sometimes not, but always with a spirit of mutual respect for ideas and opinions.

Retirement from NYU is a legal transition, not a behavioral one, unless you want it to be (and I don't). Like almost everyone else—or so it seems—I plan to consult on a freelance basis. Much, perhaps most, of my involvement will be with the aforementioned digital divide issue, with the Internet Society and with other organizations that focus on related issues, such as the Markle Foundation, the United Nations Development Program, the World Bank, USAID, the Dot Force initiative, and, indirectly, the G-8.

In addition, I've just accepted a position as Executive Director of the Global Internet Policy Initiative (GIPI), a project so new that we don't even have a website yet. Sponsored by InterNews and the Center for Democracy and Technology (which do have websites!), GIPI is designed to take the Internet to the next level in developing countries. We believe that the primary obstacles to the Internet's effective deployment and use in many developing countries are no longer technical, although the build-out of infrastructure is a daunting task requiring much time and investment. Rather, the more pressing issue will be the body of legislation, regulatory practices, and policies that define the environment in which the Internet must grow.

The Internet may be ready to expand its presence, but the legal, regulatory, and policy environments are often not ready for it in the sense that they retard its growth and use, and do not allow civil society to benefit from it. Countries are beginning to realize this, but have little effective assistance in understanding what the Internet is and what helps it to deliver.[11] GIPI hopes to close this gap by assisting constituencies in developing countries to form and examine this environment and its effect on Internet readiness, and then engage the policy process to make changes. We believe that the general principles underlying such transformations are well understood, but that their implementation is clearly country specific and requires the buy-in of many sectors of the society. Thus, our modality of operation is to hire local advocates, give them guidance through directors in the field and back them up with international legal assistance from Washington and elsewhere. We focus on implementation, on the ground, based on rough consensus, applying principles that we believe are international in scope. It's going to be a challenge and an opportunity, and I'm looking forward to it.

I remain optimistic about the future, but it is not guaranteed, so I feel some obligation to work toward ensuring that my version of it prevails. I hope that I will continue to serve on the Board of the Internet Society and as its Vice President for Education, and continue to formulate projects and try to obtain funding for them from benevolent donors. One of my aspirations for this coming year is to be able to establish, with a great deal of help from others, regional Internet training centers in Africa and in the Middle East. I'm sure that other projects that arise will maintain my average altitude over time.

I wish you all good fortune in dealing with the future of information technology at NYU. It's sure to be an exciting time, and under Marilyn McMillan's capable and perceptive leadership I'm certain that you will do well.

Thank you very much.

[11] The United States is not exempt from such confusion. After substantial argument at the court hearing in Philadelphia several years ago challenging the constitutionality of the Communications Decency Act, the judge ruled that the Internet did not correspond to any of the analogous communications methods invoked during the trial, but rather was a new medium with new properties that distinguished it from previous forms of communication.