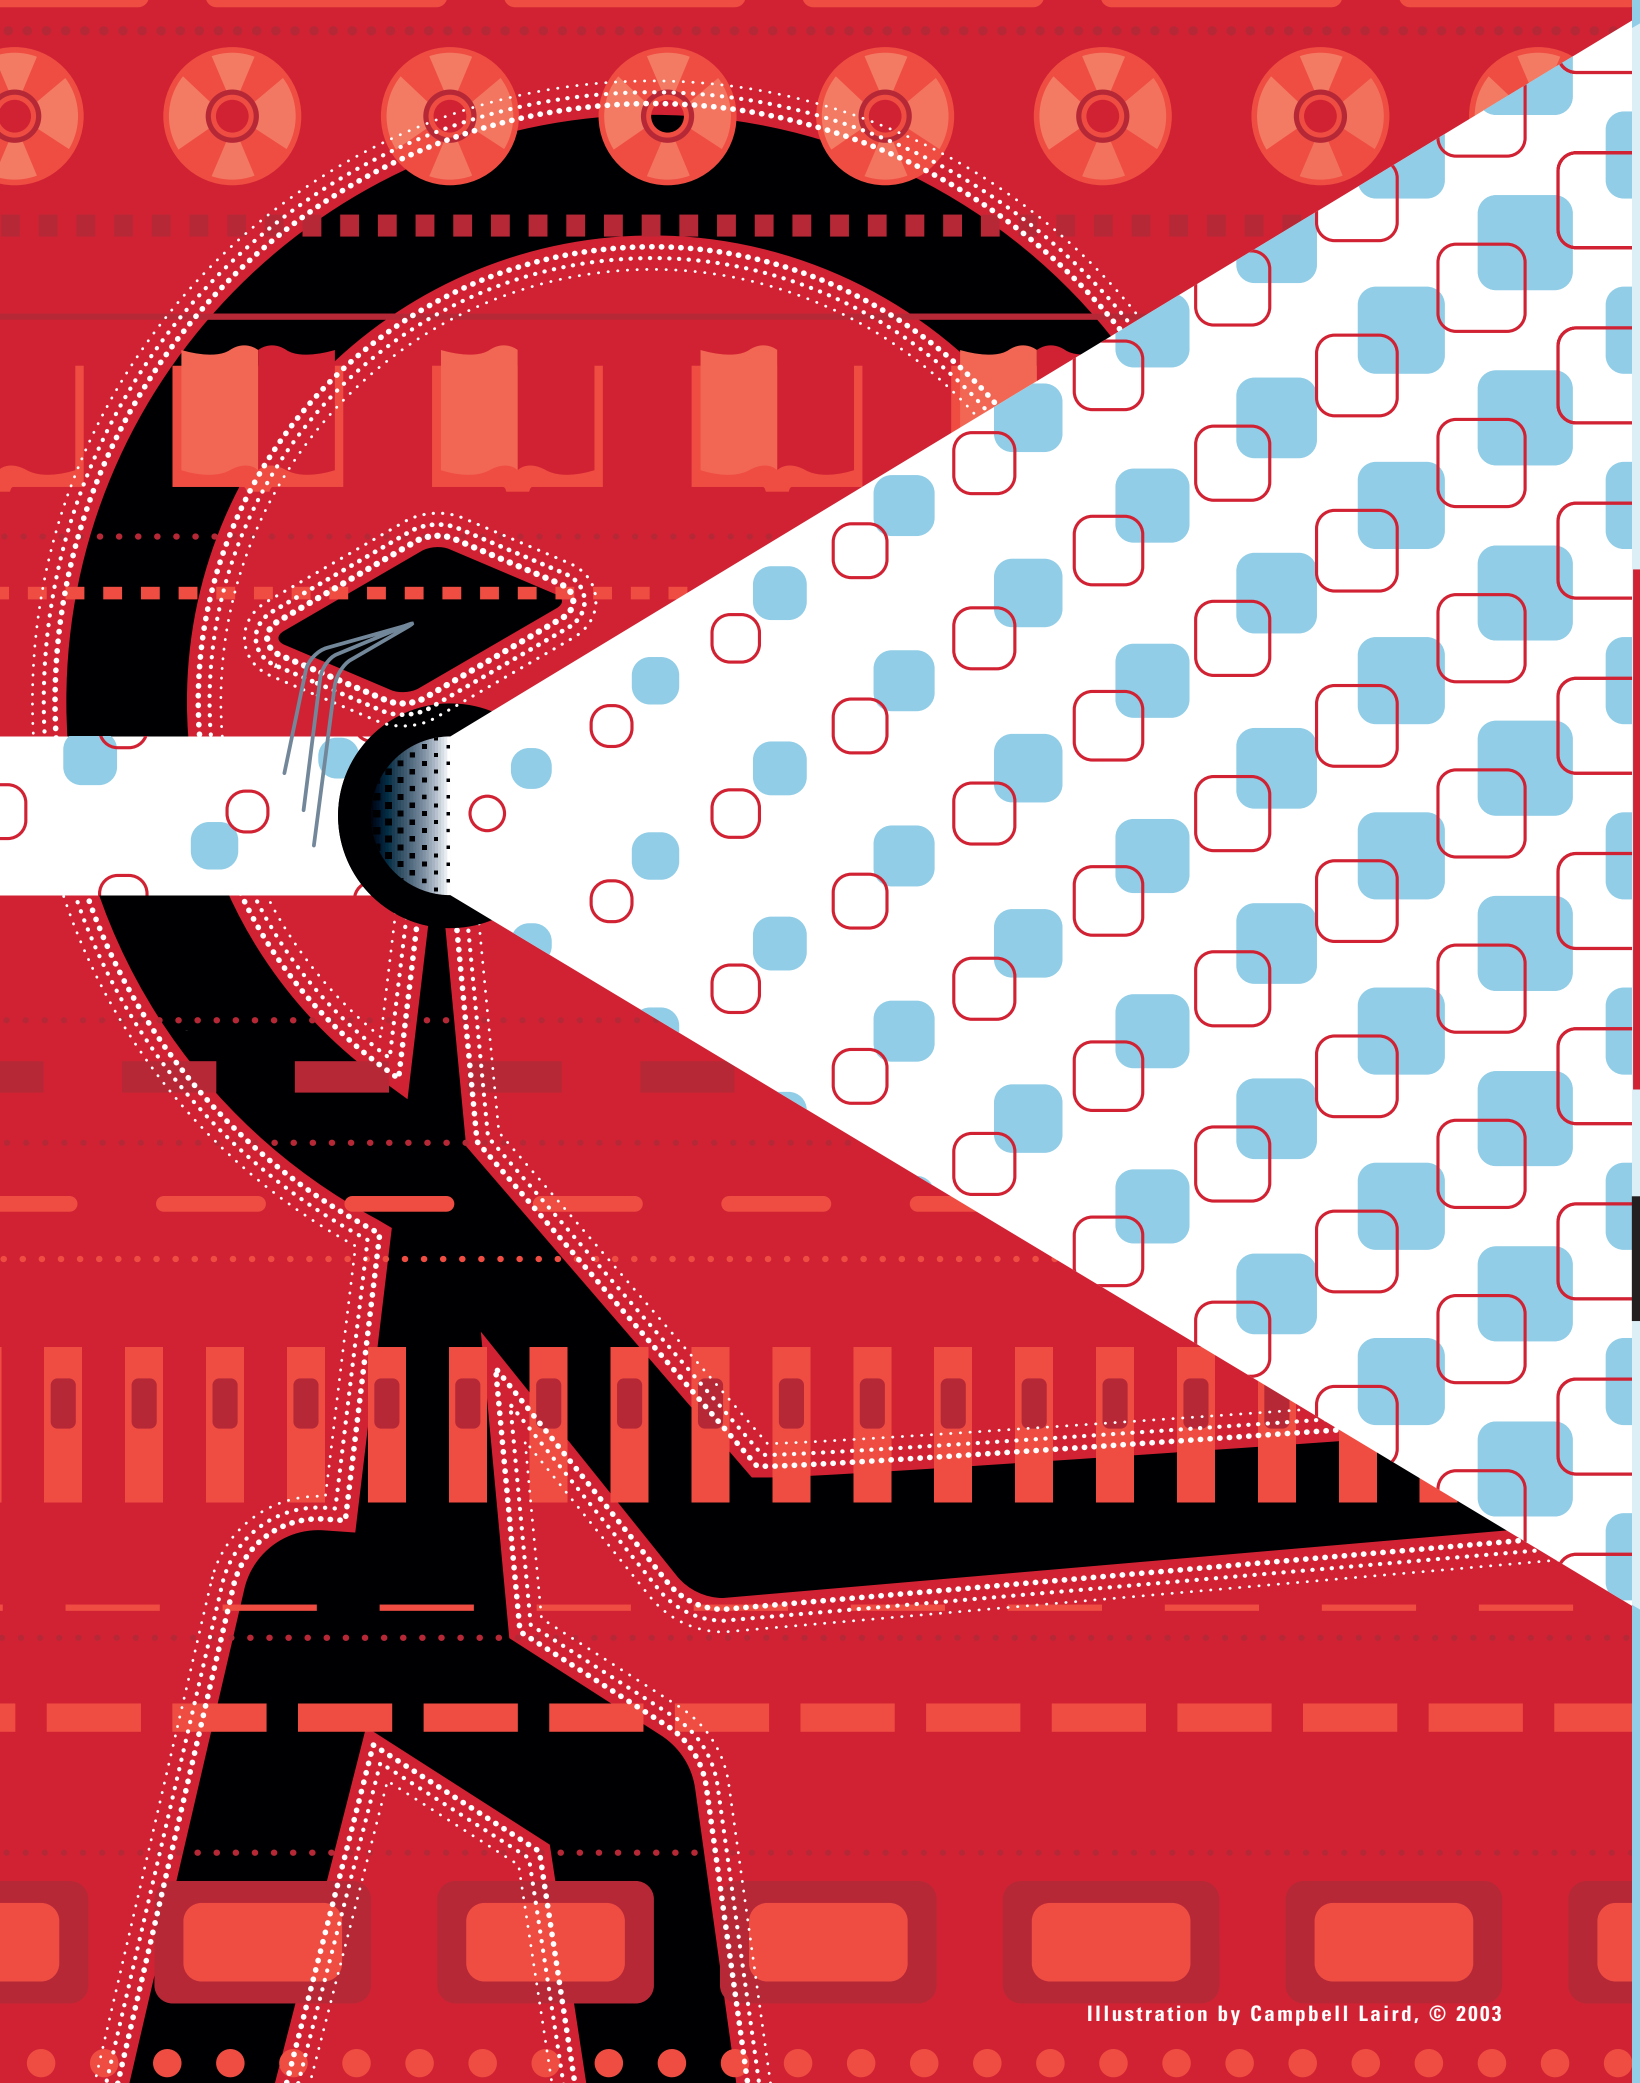

Illustration by Campbell Laird, © 2003

Beyond Bandwidth...

By James D. Bruce

Network bandwidth—which today enables the access, delivery, and exchange of information at speeds that significantly affect daily life worldwide—has gone from a national backbone of 56 Kbps (kilobits per second) in the late 1970s to 10 Gbps (gigabits per second) today. This is a factor of about 200,000 in twenty-five years. Stated more dramatically, academic and research backbone bandwidth has grown at 60 percent per year—doubling every eighteen months—for the past two and a half decades.

James D. Bruce (jdb@mit.edu) is Professor of Electrical Engineering and Vice President for Information Systems at the Massachusetts Institute of Technology.

The explosive growth of network bandwidth has played a critical role in the increased use of technology to enhance the day-to-day learning and working experiences of college and university students, faculty, and administrators. Few would disagree that computing speed and capacity, coupled with increasingly innovative applications, have created unprecedented opportunities for increased productivity, creativity, and collaboration. And the growth in network bandwidth has enabled vast changes in technology deployment on a global scale, with global consequences.

Without question, commercial network providers and higher education institutions—especially that part of the higher education community connected to Internet2—provide network bandwidth and capabilities far beyond what could only be imagined a few short years ago. Can the current pace be sustained? What are the implications and possibilities if the growth curve is maintained? What are the costs? What applications will be running five, ten, twenty-five years from now? What will the classroom look like? The research lab? The office? The higher education institution? No one knows. But trying to peek under the curtain of the future—"beyond bandwidth," if you will—is an exciting mission.

Of course, any such future-gazing involves uncertainty. In the world of computing, two well-known predictions of the future were well off the mark, yet they are very instructive nonetheless. In 1943, Thomas Watson Sr., the founder of IBM, said, "I think there's a world market for maybe five computers." By 2001, the annual worldwide production of personal computers had exceeded 135 million units.[1] In 1977, Kenneth Olsen, the founder of Digital Equipment Corporation (DEC), said, "There is no reason why anyone would want a computer in their home." Today, according to the U.S. Commerce Department's 2000 statistics, more than 51 percent of U.S. households own at least one personal computer. That's more than double the 1993 figure of 23 percent of U.S. households.[2]

That giants such as Watson and Olsen were so far off the mark speaks more to the context of computing in their day than to a lack of vision. Even the people with vision—and these were certainly men of vision—could not conceive of a world that would be so reliant on computers. Today, we live in a networked, digital world—one that is becoming even more so every day. And like Watson and Olsen, we are restricted in our vision by the context of computing today.

The Perspective of History

The research that led to the Internet began in the early 1960s with Leonard Kleinrock's work on packet-switching theory, first at MIT and later at UCLA. By 1966 Thomas Marill and Larry Roberts had written "Toward a Cooperative Network of Time-Shared Computers," which is considered the first plan for the ARPANET (ARPA is the Advanced Research Projects Agency, part of the U.S. Defense Department). In the mid-1970s, based on their work to understand whether packet switching as a technology would be useful for implementing computer-based command-and-control systems for the military, Vinton G. Cerf and Robert H. Kahn began specifying the Transmission Control Protocol (TCP).[3]

The ARPANET came into existence with four nodes in 1969—at UCLA, the Stanford Research Institute, the University of California at Santa Barbara, and the University of Utah—and reached twenty-three hosts at fifteen nodes by 1971. By 1984, one thousand hosts were on the ARPANET, and in 1985 symbolics.com was assigned the first registered domain. It wasn't until 1986 that the first commercial networking product—routers—came into being. Also in 1986, the National Science Foundation (NSF) established the NSFNET. In 1990, the ARPANET was decommissioned. In 1995, the NSFNET was retired, after less than a decade of service. In the NSFNET's nine years of existence, its bandwidth increased from 56 Kbps to T-3 (44.7 Mbps, or megabits per second) and traffic reached 10 trillion bytes per month.

Although commercial connections were first permitted on the Internet in 1989, the public did not really become aware of the Internet until around 1994, primarily as a consequence of the World Wide Web.[4] In the mid-1980s, few students had e-mail or Internet access. There were no MP3 files or instant-messaging. By contrast, freshmen in the class of 2006 most likely have computers in their rooms, full Internet access, and a cell phone in their pocket. In July 2002, the *New York Times* reported that 61 percent of U.S. adults use the Internet, up from 46 percent two years earlier. According to the *UCLA 2001 Internet Report,* these individuals spend an average of 9.8 hours per week—12 hours if they have a broadband connection—online using e-mail, instant messaging, browsing, buying, finding entertainment information, and reading news.[5] It is not an overstatement to conclude that today's students think of the Internet the way their parents and their grandparents—and even their older siblings—viewed electricity: ubiquitous, and noticeable only when unavailable.

While the world was capitalizing on the Internet, the research and education communities continued to advance network technology. Internet2 had its beginnings in 1997 and grew quickly to span the United States with a 2.4 Gbps backbone. Today, Internet2 has two hundred college and university members and over sixty corporate partners. The network also peers with research and educational networks in over forty countries around the world. Internet2's primary goal is to enable a new generation of applications and to transfer these capabilities to the global Internet.

Networking has gone from theory to reality, from an idea to a national backbone of 56 Kbps in the late 1970s to 10 Gbps today.

Today's Computing Environment

In the mid-1980s MIT, enabled by strong partnerships with DEC and IBM, launched Project Athena. The goal was to make high-performance, networked workstations available to faculty and students to drive educational change. These workstations were called "3M" machines—one million instructions per second, one megabyte of RAM, and one megapixel monochrome displays.

This environment was actually realized in the mid-1980s with the DEC MicroVAX II. At that time, MIT's campus backbone was a 10 Mbps ProNet ring, and building networks were shared 10 Mbps Ethernets. Off-campus connectivity was via the ARPANET and was about 100 Kbps, since MIT had two ARPANET

IMPs (Interface Message Processors) at that time.

Today, my Apple Titanium G4 PowerBook runs at 800 MHz. It has a one megapixel by 24-bit display, 512 MBs of RAM, and a 60-GB hard drive. My machine supports Ethernet at 10 or 100 Mbps. MIT has 322 Mbps of commodity Internet bandwidth and connects to the Northern Crossroads, the local Internet2 connection point, at 1 Gbps. Besides weighing much, much less, my PowerBook has some 20,000 times more computational performance, at less than half the price, than the MicroVAXs that MIT deployed about seventeen years ago.

This example—comparing the MicroVAX II from the early days of Project Athena to my G4 portable—points to three scaling laws:

- Moore's law—processor capabilities double every eighteen months
- Saltzer's observation—solid-state and rotating memory double every twelve or so months
- Metcalfe's law—the price of commodity bandwidth decreases by 50 percent every nine months

So in a decade, if you believe, as I do, that these three scaling rules will remain valid, a machine costing about $2,500 will run one thousand times faster and have one thousand times more memory than current machines. And every college or university will likely be able to acquire one thousand times as much external bandwidth for what they are now paying. But quite likely even that bandwidth will not be sufficient for the research, academic, and community needs of students, faculty, and staff.

Scaling rules like these reinforce the comments of economists who are coming to the conclusion—given experiences with other technologies—that we are still in the early stages of discovering the productive uses of computing and communications technologies.[6]

Leading-Edge Applications

This brings us to a point where we must look at the "beyond" in all dimensions associated with computing and communication technologies. In a decade, computing will be as different from what it is today as computing today is different from what it was in the early 1990s. How are we going to use the capabilities that we will have? What can we do with the capabilities and the technology to improve how we learn? What new research opportunities does the technology create? What impact will the technology have on the cost of education and research?

Today's students think of the Internet the way their parents and their grandparents—and even their older siblings—viewed electricity: ubiquitous, and noticeable only when unavailable.

Even with the new technological breakthroughs that will occur—like the Web or instant messaging from the last decade—the best predictors of how we will use technology in the future are the advanced applications of today. Similarly, the best predictors of the new tools and capabilities that we will need to make fuller use of our computers and networks are the ones that are inhibiting today's leading-edge applications. This remains so because the lead time from invention to full production is still measured in years.

Let's consider a few leading-edge applications—and remember, "leading edge" is in the eye of the beholder.

Colonial Williamsburg

Colonial Williamsburg is an excellent example of a living museum: period buildings, interpreters, artifacts, and opportunities to learn. In addition to being a physical place, the museum also brings this living history into classrooms across the country, transforming the way history is taught. The museum has developed several video series—for example, "Becoming Americans" and "A Day in the Life"—along with electronic fieldtrips. These video materials are typically delivered to the classroom by public broadcast, although scheduled streaming video over the Internet is also utilized. When the materials are used in the classroom, they can be accompanied by reproduction artifacts that provide individualized, hands-on experiences. Each artifact furnishes an opportunity to touch, examine, use, and understand things from another time. For example, after watching an interpreter portray a colonial lady, schoolchildren can open a "Hands-on-History" kit to see exactly what she has in her pocket: a fan, a purse with a few coins, sealing wax, and a hair curler. Certainly these materials can make learning American history substantively more exciting than in the past.

Suppose that teachers could stream these video resources into their classrooms via the Internet when they wanted to use the material rather than conforming to a rigid broadcast schedule. Suppose the materials could become an enhanced, individualized learning environment for teachers and students for personal study and for reference use. Suppose further that teachers could easily take segments from various video materials and bring together multiple presentations on the same subject for class or individual study. All of these technology-enhanced methods would further improve teaching and learning.[7]

Million-Book Digital Library

Let's turn from pictures to words. It is estimated that there are about 110 million books and documents in the world today. To demonstrate the feasibility of making all these books available online, Professor Raj Reddy of Carnegie Mellon University, with support from the NSF and the governments of India and China, is undertaking to build a million-book digital library by the year 2005.

In this project, books are scanned at 600 dpi, and the resulting page images are stored. A searchable index is being generated using Optical Character Recognition (OCR) technology to generate text. Professor Reddy estimates that some 50–60 MB will be required to store the contents of a single book and that 100 TB will be required to store one million books. Using today's technology, this translates into ten racks of equipment occupying a total of forty square feet of floor space and costing several million dollars. Using 2005 technology, this will require only two racks of equipment occupying four square feet of floor space and costing less than half a million dollars.[8]

The best predictors of the new tools and capabilities that we will need to make fuller use of our computers and networks are the ones that are inhibiting today's leading-edge applications.

Virtual Grand Rounds

One of the teaching mechanisms found in teaching hospitals is grand rounds. Typically in grand rounds, doctors present new findings about the treatment of difficult diseases, often illustrated by specific cases. This is great for medical staff who work and learn in teaching hospitals, but it represents a missed opportunity for medical personnel outside such a setting.

MIT's Center for Experimental Pharmacology and Therapeutics, a part of the Harvard-MIT Division of Health Sciences and Technology, has established the Fungal Infections Virtual Grand Rounds as a practical guide to the understanding and management of invasive fungal infections. Its Web site (http://www.figrandrounds.org) includes lectures—in video and slide formats—by distinguished faculty and instructive cases that make this material relevant to anyone engaged in the management of fungal disease.[9] A quick search of the Web provides access to thousands of similar sites worldwide.

Classrooms, Research Computing, and Grid Technology

MIT is conducting a number of projects to use computing and communications technology for learning. The Institute is working with two universities in Singapore to develop five new master's degree programs. About half the subjects in the programs are taught from MIT using streaming video over Internet2—some twelve hundred hours in forty subjects so far. Classrooms at MIT and in Singapore are equipped to originate classes, and students at MIT and the two Singapore universities are enrolled in these subjects. Educational evaluations have demonstrated that students learn equally well in the classrooms at both ends of such a learning system.[10]

MIT is also in the process of putting the content of all its classroom subjects on the Web. The OpenCourseWare (OCW) project will make available lecture notes, problem sets, supplemental readings, and other materials for the two thousand or so subjects that are offered each year. The first subjects are now available on MIT's Web site (http://ocw.mit.edu). Although the primary audience for these materials will be teachers around the world, the content will be freely available to everyone. MIT initiated this project for two basic reasons:

1. Over the years, textbooks developed at MIT have had a worldwide impact on the way many subjects have been taught. MIT—along with the Andrew W. Mellon Foundation and the Hewlett Foundation, the primary sponsors—believes that well-developed, online content will be the equivalent of the textbook in the future.
2. This project will make it easier for MIT faculty to provide their students with well-developed materials, to share their ideas and teaching materials with each other, and thus to improve educational quality at MIT.

To provide a framework for use by faculty and others developing online content such as OCW, MIT has undertaken—with support from the Mellon Foundation and in collaboration with Stanford University, Dartmouth College, North Carolina State University, the University of Michigan, the University of Pennsylvania, the University of Wisconsin at Madison, the University of Washington, and Cambridge University—the development of an open and extensible architecture for learning technology targeted to the needs of higher education. Known as the Open Knowledge Initiative, or OKI (http://

web.mit.edu/oki), this architecture specifies how the components of a learning technology environment communicate with each other and with other campus systems.

By clearly defining points of interoperability, the architecture allows the components of a complex learning environment to be developed and updated independently of each other. At the core of OKI is a set of application programming interfaces (APIs) that realize the OKI architecture. The project is providing Java versions of these APIs. The OKI partners and the developer community are providing open-source examples and reference implementations of learning technologies that make use of the APIs. MIT, Stanford, the University of Michigan, and other institutions are using this architecture to build learning-management systems for use on their campuses. Learning environments must be flexible enough to adapt to a wide range of instructional requirements and styles yet stable enough to allow faculty and students to concentrate on teaching and learning and not on the technology itself.

In addition to using information technology in the classroom, faculty, staff, and students are also heavily involved in research. Historically, faculty research often involved a single faculty member, maybe one or two research staff, and a small number of graduate and undergraduate students. The exact composition of the group depended on the field of work and the culture of the particular college or university. Today, in the era of "big science," many projects involve not just a few but often hundreds or even thousands of researchers. These researchers and their research instruments are geographically distributed around the globe.

One example is the Gemini Observatory, which consists of twin 8-meter telescopes, one located on Mauna Kea in Hawaii and the other on Cerro Pachón in the Chilean Andes. Internet2 links the observatory's two telescopes. An astronomer in either location can have a real-time "presence" in the other control room as the observing program is executed on both telescopes. With sufficient bandwidth, an astronomer can extend the real-time presence to his or her remote office or laboratory as an observing program is run. As Dr. C. Matt Mountain, the director of Gemini Observatory, remarked at the dedication of Gemini South, "Welcome to the era of point-and-click astronomy."[11]

In addition, a phenomenon called "grid computing" is developing around the globe. Fundamentally, grid computing refers to a collection of resources—computing power, storage, instruments, sensors, and processes for accessing the resources—organized in an environment that enables resource sharing and coordinated problem-solving in dynamic, multi-institutional collaborations. An example is the TeraGrid, a multi-year effort to build and deploy the world's largest, fastest, and most comprehensive distributed infrastructure for open scientific research. This transcontinental supercomputer will have clusters of high-end microcomputers at four different institutions; together, the microcomputers will deliver about 14 trillion floating-point operations per second (teraflops). The machine will also have 650 TB (terabytes) of networked storage and be linked via a 40 Gbps network. The resulting virtual computer will enable scientists to tackle some of the most computationally challenging tasks on their research agendas.[12]

A final example of research computing and grid technology is the Large Hadron Collider, an experimental high-energy physics project. The physical instrument is located at CERN, the European Organization for Nuclear Research, and will generate a data stream of 1 to 100 Gbps when it is operating. Most of this data is stored on tape in libraries, which are expected to increase by 5 to 8 PBs (petabytes) per year, equivalent to some 10 million CDs annually—enough to fill about 20 miles of shelving. When a research collaborator wants to analyze data from one of the several collider experiments, the researcher must stage the data to disks local to the computational resources that will be used. This will involve replicating the data, cataloging it, and moving these very large data sets across the global Internet. (Moving 1 PB at 10 Gbps takes about 15 minutes using the full network. Moving 1 PB across today's campus networks would take much longer—250 hours at a 10 Mbps bandwidth for the entire path.)

The Next Challenges

Throughout my review of these examples, I have assumed that the Web mediates between the user and the technology, vastly simplifying the interface between the two. I believe that this will become even more so in the future, with the Web becoming the de facto operating system of choice, providing access to systems and resources. But the above examples, which illustrate some of the uses being made of computing and communications technology, also reveal several challenges that we will have to overcome in order to fully reach our goals.

End-to-End Bandwidth

If Metcalfe is correct, the price of bandwidth should become so low that we will be able to acquire all the bandwidth we need. For network backbones—for example, the Internets—this seems to be the case, but for the entire path, it's not quite that simple. Campuses and communities face significant difficulties in delivering high bandwidth across that "last mile" because the infrastructure is inadequate.

At the end of 2001, higher-speed Internet access—for example, DSL or cable modems—was available to approximately 80 percent of U.S. households, yet less than 10 percent of these households were connected. Essentially, no households are connected at the bandwidths that will be needed in the near future. Recognizing the importance of broadband access in "every corner of the country," U.S. Senator Joseph Lieberman of Connecticut introduced legislation (S.2582) that calls for the development of a coherent, cross-agency strategy to facilitate deployment of 10 to 100 Mbps across all parts of the country.[13]

Not too long ago, I asked several of my colleagues, the CIOs of the IVY+ institutions, about their network infrastructure. I found that we currently deliver essentially no Gigabit Ethernet (GigE) service,

In the ideal world, network users would have a tool that could tell a user where a problem is, what type of problem it is, and the person to contact for resolution of the problem.

even though about 45 percent of our installed wiring from telecommunication rooms to desktops will support GigE. Today, about half of our desktops are supported at 100 Mbps Ethernet and the remainder at 10 Mbps, with significant variation from campus to campus. Yet, we all see the need to move rapidly to higher bandwidths, driven by the applications that we see on the horizon. Thus, we are struggling with delivering higher-speed access whether we are talking about residences and businesses or college and university desktops. In the context of delivering end-to-end services, the issue is not just the bandwidth from the local "telecommunication room" to the desktop but the bandwidth across the entire path to the other "end" from which the service is being delivered.[14]

Internet2 has formed a working group to assist campuses in their efforts to improve end-to-end performance. Its charter focuses on improving overall performance and on detecting and resolving problems wherever they occur in the networking infrastructure. In the ideal world, network users would have a tool that could tell a user where a problem is, what type of problem it is, and the person to contact for resolution of the problem. To assist in this endeavor, performance data between key nodes in the web of campuses, regional networks, gigapops, and the Abilene backbone is regularly being collected.

Adequate Network Addresses

In his keynote address at the University of Michigan's 1998 Marshall Symposium, the Internet pioneer Vinton G. Cerf remarked: "It does seem, as time goes on, that more and more things are going to be connected to the Internet.... By 2005, the light sockets will be on the Internet, and when we're controlling the lights, we won't be turning the current off and on. We'll be sending SMNP Internet control packets to the light sockets, telling them to turn the light off and on and get brighter or dim."[15]

Whether or not lightbulbs will be controlled by the Internet, obtaining adequate Internet addresses is difficult enough today, particularly outside the United States. The lack of addresses is the result of the Internet Protocol (IP) suite known as IPv4, which uses 32-bit addresses. This results in about 4 billion addresses, most of which are assigned to the United States. Among the handful of U.S. organizations that have Class A addresses, each has about as many network addresses as the entire country of China.

The solution to this problem is found in a new protocol suite: IPv6, which uses 128-bit addresses, essentially equivalent to an address for each grain of sand on the planet. Internet2 can already route network traffic using IPv6, and Japan has mandated IPv6 on government and educational networks by 2005. In addition to providing the address space that the Internet needs, IPv6 will enable quality-of-service guarantees that will improve IP telephony, video-conferencing, multicasting, and peer-to-peer applications.

Deployment of IPv6 more broadly awaits the incorporation of the protocol into the operating systems of desktop products. This will begin to happen in the next releases but will not reach a peak for another two years. Accomplishing full deployment will require readdressing all of the hardware that is in place and will likely be accompanied by significant confusion. (As a comparison, think about the confusion that would result from renaming and renumbering all the streets in a city.)[16]

Middleware

As more and more applications are built to provide access across organizational boundaries, mechanisms and policies are needed for secure interinstitutional authentication and authorization. These mechanisms are generally called "middleware." Middleware is concerned with managing resources so that they work in an integrated fashion and the user does not have to make separate arrangements to locate or use remote resources, share data, or ensure levels of security.

Many campuses have developed intra-institutional identity, authentication, authorization, and directory services to meet some of these needs. Higher education institutions will have to extend what is already in place to more fully deliver services to meet the needs of the entire campus. And institutions will have to establish a full set of services to address the growing number of interinstitutional collaboration arrangements that are developing.

Internet2 has taken a leadership role in developing middleware standards and toolkits and also in initiating pilot deployments.[17]

Security

A decade ago, MIT experienced about two to three attempts to break into its machines each year; in 2001, the number was 5,600—and the number is growing rapidly. A new computer placed on the MIT network will be scanned within a matter of minutes, and poorly configured Windows machines will typically be compromised within fifteen minutes. At MIT, as on most other campuses, many departmental or specialized computer systems are managed by students or research staff. Unfortunately, for many of these individuals, security is an irritating side issue.[18]

A recent EDUCAUSE/Internet2 paper urged higher education institutions to address three goals:

1. Detect and prevent attacks against campus systems
2. Detect and prevent attacks going to off-campus computer systems
3. Secure vital campus systems and data against attack[19]

Reaching these goals will require significant work—on the part of both the campus information technology organization and the campus user community, as well as on the part of operating system and application vendors. There are no simple answers.

Going "Beyond"

As I tried to think of a way to pull this commentary together, I kept coming back to my own experiences in computing. I first used a computer in 1958. It used punched cards for input and output, and its memory consisted of a two-thousand-word rotating drum. Running my "job" often took sixteen hours. The laptop I have today is much more powerful than all of the computers that existed in the world at that time—combined.

In 1958, almost all computing applications I knew about were either financial calculations of some sort or work on engineering design or models. Today, computing enters our lives at almost every point of our existence, affecting us in many ways. For example, last winter I learned that Medical Media Systems Inc. (http://www.medicalmedia.com/) had developed the capability to take a full-body CAT scan and, using advanced 3-D imaging techniques, create a highly accurate, patient-specific, 3-D rotatable anatomical model. Such models can provide detailed visualization for surgeons in their planning of complex surgical procedures. This technology was of particular interest to me because my eighteen-month-old granddaughter, Elizabeth Joy, was about to undergo surgery to remove a very large neuroblastoma that had been shrunk by chemotherapy. The imagery showed that the tumor was well isolated, a situation that was confirmed by the surgery performed to remove it. Today, after more chemotherapy and radiation therapy, Lizzy runs and plays like any other normal two-year-old and is routinely monitored to detect any reoccurrence of the tumor.

This and the other examples I have shared point to a rich fabric of use and provide glimpses of what the future will hold as we go "beyond bandwidth." The barriers I've mentioned will all succumb, in one fashion or another, and we will live in a computing world in which there will always be enough bandwidth. New applications of computing will come on the scene and will "knock our socks off"—in the same way that e-mail and the Web did when they first appeared. Yet new operational challenges will arise as well, of course. The future is here, but we still have work to do. Our challenge is to go "beyond bandwidth"—to face this future openly and to realize the promise it holds.

Notes

1. Watson's comment is quoted in BBC News, "50 Years of Computers," June 18, 2001, <http://news.bbc.co.uk/1/hi/in_depth/sci_tech/2000/dot_life/1394413.stm> (accessed November 6, 2002); for the 2001 figure, see Plunkett Research Ltd., "Trends in Personal Computers and Internet Access" <http://www.plunkettresearch.com/technology/ecommerce_computer_trends.htm> (accessed November 6, 2002).
2. Olsen's statement is quoted in S. L. Mintz, "Getting Real," *CFO Magazine,* November 1, 1999, <http://www.cfo.com/printarticle/0,5317,1388,00.html> (accessed November 6, 2002); for the statistics on U.S. households with computers, see U.S. Census Bureau, U.S. Department of Commerce, "Home Computers and Internet Use in the United States: August 2000," September 2001, <http;//www.census.gov/prod/2001pubs/p23-207.pdf> (accessed November 6, 2002).
3. Thomas Marill and Lawrence G. Roberts, "Toward a Cooperative Network of Time-Shared Computers," fall AFIPS Conference, November 1966, <http://www.packet.cc/files/toward-coop-net.html> (accessed November 6, 2002); Robert H. Kahn and Vinton G. Cerf, "A Protocol for Packet Network Interconnection," *IEEE Transactions on Communications,* May 1974, <http://www.worldcom.com/global/resources/cerfs_up/technical_writings/protocol_paper/> (accessed November 6, 2002). In 2001, Kleinrock, Roberts, Kahn, and Cerf were all honored by the National Academy of Engineers as recipients of the Charles Stark Draper Prize for their collaborative work that led to the invention of the Internet.
4. Sources for the Internet timeline data include Nerds 2.0.1 Internet Timeline: <http://www.pbs.org/opb/nerds2.0.1/timeline/70s.html> (accessed November 6, 2002).
5. Technology Column, *New York Times,* July 22, 2002; UCLA Center for Communication Policy, "Surveying the Digital Future," *UCLA Internet Report 2001.*
6. "The Economy," *Fortune,* September 9, 2002, 92.
7. On Colonial Williamsburg, see Richard L. McCluney, "The Freshest Advices: A Little Revolution in American (History) Education," *Journal of Colonial Williamsburg Foundation,* <http://www.history.org/foundation/journal/index.cfm>.
8. Raj Reddy, "Computing, the Next Ten Years: Infinite Memory and Bandwidth—Implications for Universal Access to Information," talk presented at Georgia Tech Tenth Anniversary Convocation, April 6, 2001.
9. "'Virtual Grand Rounds': Rx for Fungal Infections," *MIT Tech Talk,* July 17, 2002.
10. "Distance Education Infrastructure and Support for SMA" (Singapore-MIT Alliance), presentation at MIT, July 2002.
11. Gemini Observatory newsletter, no. 24, June 2002.
12. On grid computing, see Patrick Dreher, "Research Computing and Grid Technology: How These Developments May Impact Higher Education IT Services," presentation at the Common Solutions Group Meeting, Seattle, Washington, September 2002; M. Mitchell Waldorp, "Grid Computing Could Put the Planet's Information Processing Power on Tap," *Technology Review,* May 2002; and Cade Metz, "Grids," *PC Magazine,* October 1, 2002.
13. "Lieberman Calls for Adding Broadband Policy to Agenda at White House Tech Forum," press release, June 12, 2002, <http://www.senate.gov/~lieberman/press/02/06/2002612819.html> (accessed November 6, 2002).
14. On end-to-end bandwidth, see Information Technology Association of America, "Building a Positive Broadband Agenda," white paper, October 2001; and Martin Cooper, "What Ails Broadband?" *IEEE Spectrum,* September 2002.
15. Vinton G. Cerf, Marshall Symposium keynote address, Rackham School of Graduate Studies, University of Michigan, May 29, 1998, <http://www.si.umich.edu/marshall/docs/p106.htm> (accessed November 6, 2002).
16. On IPv6, see Jessica Schneider David, "What Is Ipv6?" *iQ Magazine,* July/August 2002; "Ipv6 Stands to Level the Playing Field," *InfoWorld,* July 29, 2002; Internet2, "Under the Hood" and "Ipv6 Addressing," Ipv6 Workshop, Research Triangle Park, North Carolina, March 5–7, 2002.
17. On middleware, see Alan Blatecky, Ann West, and Mary Spada, "Middleware: The New Frontier," *EDUCAUSE Review,* July/August 2002; and Ken Klingenstein, "Middleware 101," <http://middleware.internet2.edu/overview/middleware-101.pdf>, and "PKI: News from the Front and Views from the Back," <http://middleware.internet2.edu/overview/PKI-news-and-views.pdf> (both accessed November 6, 2002).
18. Charles M. Vest, "Openness, Opportunity, and Securities in Universities: A National Challenge," address to the National Association of College and University Attorneys, Boston, June 2002; Bob Mahoney, "The Problem of Security in a Large University Network," presentation, April 17, 2002, MIT.
19. EDUCAUSE/Internet2 Computer and Network Security Task Force, "Higher Education Contribution to National Strategy to Secure Cyberspace," (draft), June 26, 2001, <http://www.educause.edu/security/cybersecurity_draft.doc> (accessed November 6, 2002).