VOLUME 30, NUMBER 2
SUMMER 2000

The BRIDGE

The
BRIDGE

The Bridge (USPS 551-240) is published quarterly by the National Academy of Engineering, 2101 Constitution Avenue, N.W., Washington, DC 20418. Periodicals postage paid at Washington, D.C.

Vol. 30, No. 2, Summer 2000

Editor: Karla J. Weeks

Production Assistants: Penelope Gibbs, Kimberly West

Postmaster—Send address changes to *The Bridge*, 2101 Constitution Avenue, N.W., Washington, DC 20418.

Papers are presented in *The Bridge* on the basis of general interest and timeliness in connection with issues associated with engineering. They reflect the views of the authors and do not necessarily represent the position of the National Academy of Engineering.

The Bridge is printed on recycled paper. ♻

A complete copy of each issue of *The Bridge* is available online at http://www.nae.edu/TheBridge in PDF format. Some of the articles in this issue are also available as HTML documents and may contain links to related sources of information, multimedia files, or other content.

THE NATIONAL ACADEMIES

National Academy of Sciences
National Academy of Engineering
Institute of Medicine
National Research Council

The **National Academy of Sciences** is a private, nonprofit, self-perpetuating society of distinguished scholars engaged in scientific and engineering research, dedicated to the furtherance of science and technology and to their use for the general welfare. Upon the authority of the charter granted to it by the Congress in 1863, the Academy has a mandate that requires it to advise the federal government on scientific and technical matters. Dr. Bruce M. Alberts is president of the National Academy of Sciences.

The **National Academy of Engineering** was established in 1964, under the charter of the National Academy of Sciences, as a parallel organization of outstanding engineers. It is autonomous in its administration and in the selection of its members, sharing with the National Academy of Sciences the responsibility for advising the federal government. The National Academy of Engineering also sponsors engineering programs aimed at meeting national needs, encourages education and research, and recognizes the superior achievements of engineers. Dr. William A. Wulf is president of the National Academy of Engineering.

The **Institute of Medicine** was established in 1970 by the National Academy of Sciences to secure the services of eminent members of appropriate professions in the examination of policy matters pertaining to the health of the public. The Institute acts under the responsibility given to the National Academy of Sciences by its congressional charter to be an adviser to the federal government and, upon its own initiative, to identify issues of medical care, research, and education. Dr. Kenneth I. Shine is president of the Institute of Medicine.

The **National Research Council** was organized by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purposes of furthering knowledge and advising the federal government. Functioning in accordance with general policies determined by the Academy, the Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in providing services to the government, the public, and the scientific and engineering communities. The Council is administered jointly by both Academies and the Institute of Medicine. Dr. Bruce M. Alberts and Dr. William A. Wulf are chairman and vice chairman, respectively, of the National Research Council.

Editorial

David D. Clark

Marjory S. Blumenthal

Meeting the Information Society's Needs

Information technology has become so pervasive, at least in developed countries, that it is the stuff of jokes and cliches. More people are buying and using information technology in more ways, more people depend on it for business and personal activities, and at the same time more people are coming to understand the limitations of the state of the art. Even though this technology is becoming pervasive, it is only in its adolescence. The potential of information technology has only begun to be mined.

The Computer Science and Telecommunications Board (CSTB) is the unit of the National Academies that assesses how information technology and its uses are evolving, and how information technology influences and is influenced by economic, social, legal, and public policy factors. The work of the board is guided by a group of 20 leaders with expertise that spans the range of computer science and telecommunications, as well as complementary fields reflecting the broader impact of information technology.

The articles in this volume draw from CSTB's repertory and illustrate its breadth: David Messerschmitt was the co-chair of a major project on information technology research and development; Randy Davis chaired a major study on intellectual property in the networked world; and Steven Bellovin was a key committee member in a recent project on the trustworthiness of networked information systems. The reports of these projects have been widely circulated, helping to educate people in a range of government, industry, and academic contexts. Each addresses where and how information technology needs to be improved because of the context in which it is used. Each addresses the larger policy environment in which information technology exists.

CSTB's work illuminates a number of themes that relate to the future of information technology. Of course, the rise of the Internet is one of them. CSTB has been producing a series of reports that describe how the Internet works, how it is being used, and what it implies for the economy as a whole. Another major theme is the trustworthiness of information technology—security, reliability, and protections for privacy and other concerns of users. Future topics that CSTB is planning to examine include privacy in the online world, the evolving information ecology, economic transformations deriving from information technology, and the collision between public safety and civil liberties concerns associated with the Internet.

The adolescence of information technology is a time of confidence. We know that a lot of ideas can and do work in practice; we know that information technology can become truly pervasive. But with that knowledge of what is possible comes recognition that there are choices to be made. What kind of technology should be developed, and whose judgment or preferences will carry the day?

Increasingly, those choices will be shaped by the social and economic impacts of information technology and the scrutiny of government entities. There is more talk about regulation, and the incentives are changing, not only for the businesses that design and market information technology products, but also for the research community at large. As Messerschmitt discusses, how to make information technology for social applications is a huge challenge; as Davis discusses, business models and law interact with the development

---

David D. Clark, NAE, is a senior research scientist for the MIT Laboratory for Computer Science and chair of the Computer Science and Telecommunications Board of the National Academies. Marjory S. Blumenthal is executive director of the Computer Science and Telecommunications Board.

of technical mechanisms to protect intellectual property on the Internet; and as Bellovin discusses, the rise of criminal as well as benign uses of the Internet stimulates law enforcement interest that can affect the design and operation of the Internet.

We hope the articles in this volume stimulate your interest in our work, and we welcome suggestions of topics that we can address in the future to maximize the potential of information technology.

David D. Clark

Marjory S. Blumenthal

Getting the Most Out of Information Technology

David G. Messerschmitt and Jerry R. Sheehan

Information technology research must address problems of large-scale systems and social applications if society is to get the most out of the technology.

David G. Messerschmitt

Jerry R. Sheehan

Continued advances in information technology (IT) have resulted in dramatic transformations in many important societal functions, including commerce, education, health care, and delivery of government services. Networking in particular has enabled the development of increasingly powerful applications that span multiple organizations, linking people and information in unexpected ways. Nevertheless, IT systems have failed to live up to their full potential in many applications—crashing unexpectedly, suffering security breaches, producing unanticipated results, and failing to scale up to meet growing demand or critical requirements. These

---

David G. Messerschmitt, NAE, is Roger A. Strauch Professor of Electrical Engineering and Computer Sciences, University of California, Berkeley. He co-chaired the NRC committee that produced the report Making IT Better, *upon which this article is based. Jerry R. Sheehan is a senior program officer for the NRC's Computer Science and Telecommunications Board. He served as study director for* Making IT Better.[1]

problems underscore the shortfall in our understanding of many issues related to large-scale IT systems and their applications. They suggest that achieving the promise of IT will require new research efforts. Drawing upon a recent report from the National Research Council (NRC) Computer Science and Telecommunications Board on the future of IT research (2000), this article argues for an expansion of the IT research agenda to address these increasingly important problems and calls for a broader set of constituencies to fund, participate in, and perform IT research to build the base of understanding needed to get the most out of IT.

Companies experience chronic difficulties in developing and deploying new IT systems.

The movement of IT out of the laboratory and the back rooms of large organizations over the past few decades has allowed computing and communications systems to touch many aspects of personal and work life. Due in large part to the Internet, people routinely use IT to communicate with friends and relatives, plan trips, shop, and enjoy entertainment. Businesses use IT to buy and sell goods and services, keep track of payroll and inventory, facilitate collaborative work, and control production lines. Government agencies use IT to provide benefits and information to citizens, process taxes and grant applications, and procure supplies. A growing number of critical infrastructures, from the nation's power grid and air traffic control systems to its financial and health care systems, increasingly benefit from (and have become more reliant on) IT. Many consumer products, such as automobiles and refrigerators, integrally incorporate IT to control their operation.

The economic impact of these IT-based transformations is significant. The United States is preeminent in both the supply and the effective use of IT, and many believe that because of IT the United States is experiencing higher economic growth and productivity gains that can be sustained well into the future. The U.S. Department of Commerce estimates that one-third of the nation's economic growth since 1992 is attributable to firms that offer IT-related goods and services, and that half or more of the nation's productivity growth in the last half of the decade resulted from the production and use of IT.[2] Some 7.4 million Americans worked in IT-producing industries or in IT-related positions in other industries in 1998, at wages significantly above the national average. All signs point toward a growing need for skilled IT workers, perhaps exceeding supply.[3]

In spite of these remarkable achievements, IT systems continue to garner a number of vexing problems. Large companies and federal agencies experience chronic difficulties in developing, deploying, and operating new IT systems. Recent surveys indicate that only one-quarter of all large-scale system development efforts are completed on time and within budget, and almost 30 percent are abandoned because requirements cannot be met (Johnson, 1999; Standish Group, 1995). Once deployed, many systems lack adequate reliability and security measures to prevent the kinds of problems that afflicted several online stock trading sites and Washington, D.C.'s Metrorail system in the past year (Layton, 1999; Meehan, 2000). Reliability and security are becoming more challenging as the Internet allows the deployment of more tightly interlinked and interdependent IT systems.

Beyond these obvious failures are missed opportunities resulting from either the failure to apply IT where it might benefit or the failure to apply IT most effectively. Many IT systems are not integrated well into the activities they support, undermining attempts to boost productivity. Other applications are inflexible, failing to accommodate changing needs or add new capabilities. IT often stifles or complicates changes in organizations, processes, and products, rather than enabling and simplifying them, as should be the case. Increasing intellectual resources are applied to circumventing or bypassing existing, outdated systems, obviating fresh opportunities. The most demanding applications sometimes have difficulty scaling to rapidly growing user bases.

Large-scale Systems

The causes of these problems are many and varied. Certainly, poor project management or inadequate worker skills and training contribute to many failed development efforts or operational failures. More significant is a lack of fundamental understanding underpinning some aspects of IT and its application. The design of large-scale systems—IT systems with large numbers of heterogeneous elements that interact in complex ways—is still as much art as science, and we have inadequate tools for predicting or analyzing performance. We have no set of design principles that ensure robust, scalable, and flexible large-scale systems. While designers have time-tested principles such as abstraction, layering, and modularity that make complicated systems easier to understand, more powerful approaches are needed as IT systems grow ever larger in scale, more complex, and more dynamic.

While these problems are not new—systems issues have long plagued IT as well as other engineering products—they take on added significance as the scale of IT systems increases and the scope of applications broadens. Future applications will be even larger and more complex, if only because they will be more highly distributed and more reliant on networking. The fact that such problems are easily recognizable today—and responsible for some of the development failures mentioned above—suggests that IT developments will be increasingly constrained without substantial and sustained improvements in fundamental understanding and tools.

Social Applications

Another complication is the increasingly tight coupling between IT systems and the processes they support. IT applications supporting organizational missions and business processes—in industry, government, education, health care, and other areas—are increasingly specialized to fit specific organizational missions. For example, an application supporting concurrent automobile design, with work distributed across several continents, will have some generic elements, such as information access, but also many functions specific to automobiles and the way they are designed. As specialized applications become more tightly intertwined with organizational missions and processes, IT becomes embedded in a human organization (much as an engine controller is embedded in an automobile), raising a number of questions about what work is delegated to people, what is delegated to machines, how the people are organized, how the machines are organized, and how these human and technical organizations relate to one another. A similar phenomenon occurred in the industrial revolution with the organization of manufacturing processes, but such issues now expand to all arenas of human activity. Dare we say we are experiencing the information revolution, and not just the information age?

In social applications of IT, such as concurrent design—applications that involve integral collections of people, organizational structures, and IT systems—the maximum value of IT (by productivity and effectiveness measures) cannot be tapped by simply automating existing processes. Rather, it is important to reconsider all aspects of organizational functions and processes in light of the capabilities of IT. It also becomes appropriate to reexamine IT itself, to see how it can become a more effective resource. From the perspective of research and design, social applications target sociotechnical systems; that is, systems composed of human organizations and networked computing. The IT part is itself a large and complex system, providing additional stimulus for the research into large-scale systems mentioned earlier.

The research community has historically devoted less attention to the challenges of large-scale systems and social applications than to the underlying core technologies. While the nation's research enterprise continues to thrive—IT companies allocated more than $50 billion to research and development in 1998, and government contributed about $2 billion in research funding—much of this research has focused on the development of new and improved software infrastructure and component technologies (e.g., microprocessors, memory, and communications devices) that underlie modern IT.[4] This research has produced important improvements in core processing, storage, and communications technologies, as well as in the understanding of massive computations and the management of large data sets. As a result, we have the building blocks with which to develop a broad array of

We have no set of design principles that ensure robust, scalable, and flexible large-scale systems.

systems. Nevertheless, the IT research base needs to expand more fully into systems and social applications if continued improvements in these building blocks are to be exploited effectively.

Social applications are an important target for research. Although they are increasingly pervasive—even universal—they are poorly understood. There are many opportunities for innovative new ways to apply IT more effectively, to design work processes and organizational structures, and to structure technology and tools to more effectively support applications and mitigate design challenges. Such research could also point the way to new opportunities for IT infrastructure that aids and supports applications.

Too often, IT is used to automate rather than transform processes.

Unfortunately, these opportunities for applications research are underappreciated in academic, industrial, and government research organizations. Too often, IT is used to automate rather than transform processes, or it is assumed that IT is a fixed point that can be incorporated into a social application but not changed in any way. Typical of this is the current rage in incorporating the Web into business functions without reflecting on whether the Web is in fact the best foundation.

The interactions between IT and applications should flow in two directions. IT should cause a reexamination of application domains with the goal of applying it most effectively, and applications should point to opportunities to improve IT. After all, IT rarely suffers from fundamental physical limitations—it is what we make of it, with few constraints. Most limits or shortcomings are conceptual, the result of restricted understanding. Today's IT is surely not the best it can be.

Research into social applications can take many forms. For example, it can focus on generic requirements that arise in a number of application domains, such as achieving flexibility to meet changing needs and requirements. Another generic requirement is personal privacy—one of the grand challenges of the information age. How can legitimate functions be realized (with assured user control over the dissemination of information) without a complicated patchwork of laws and regulations? Surely, it is feasible to achieve this through improvements in IT technology, at least to a greater extent than we see today.

Research can also focus on specific application domains. One goal is to positively affect the particular domain through the effective application of IT. Such research can also uncover important relationships among generic issues, for example, how flexibility interacts with privacy. Specific application domains can also be a driver, pointing to the most stringent requirements.

Conducting effective research into social applications requires the participation of a wider base of researchers. In particular, it must involve people and organizations that are major end users of IT, because their knowledge is integral to understanding the functions to be achieved, the application requirements and how they may be modified, and any constraints on organizations and processes. For example, what does flexibility really mean? Does it mean the ability to make two applications work together, or to allow two organizations to merge their systems, or to allow an application to be dramatically reshaped? What does privacy mean in different circumstances, and what forms of end-user control over information dissemination are required? Considering a range of application domains can give insight to these and similar questions. Such research also needs to include technology specialists, because only they fully appreciate the range of options that IT can provide. End users and technologists must collaborate, because the best opportunity is to mold the technical and human/organizational aspects of the sociotechnical systems together.

Encouraging Interdisciplinary Research

Other academic disciplines should participate as well. All the engineering disciplines, as well as other professional disciplines such as business and medicine, represent opportunities for social applications. The social sciences and law have many insights to offer and many opportunities to address challenging issues surrounding social applications. Social application research is not only interdisciplinary; it also requires a prolonged and deep intermingling that transcends the simple application of one discipline to another.

Encouraging the right kinds of interdisciplinary research to make progress on problems of large-scale systems and social applications requires a diversity of organizational structures and approaches. Work is needed to develop the theoretical underpinnings of systems and new architectural approaches to design and software development processes, as well as on the range of interactions between IT and social applications. Some of this work will need to be conducted by individual investigators; other work will require small interdisciplinary research teams or larger research centers that can launch broader, longer-term collaborations. Researchers should be allowed to identify and structure mechanisms best suited to the context and goals.

Unfortunately, there are obstacles to interdisciplinary research. There are cultural differences in the way researchers identify and define problems, in their experimental approaches, in their language and terminology, and in their expectations of outcomes. Administrative structures that reinforce disciplinary boundaries also get in the way. Within universities, in particular, the disciplinary nature of academic departments and research centers results in review and promotion practices that sometimes undervalue work across disciplinary boundaries. These impediments can be especially discouraging to junior faculty concerned about long-term career prospects and tenure. Although these issues are often of less concern in industrial research laboratories, which often organize research around problem areas rather than disciplines, they can nevertheless arise in less apparent ways.

Several steps can be taken to address these concerns. One approach is for universities (or industrial organizations) to establish new research and teaching units specifically targeted at interdisciplinary IT research, as is happening at some forward-looking institutions. Such schools, departments, and research centers serve not only to recognize and legitimize interdisciplinary research, but also to enable the establishment of review and tenure criteria that are aligned with the interdisciplinary area of inquiry. Another approach is to help interdisciplinary research flourish within existing academic departments by developing and promulgating review and evaluation criteria that explicitly recognize and encourage work across disciplinary boundaries without sacrificing quality and rigor.

Additional steps will be needed to help engage end users more effectively in IT research, but numerous issues need to be addressed. Should, for example, end-user organizations set up research organizations to address pertinent social applications? Perhaps so, but they should probably do this in a measured way, since they have little experience in managing such research and may have difficulty in attracting top IT research talent.

Conducting research into social applications requires a wider base of researchers.

More promising for the near term are mechanisms that directly involve end-user organizations in IT research but retain the actual performance of the research in existing industrial, government, and, especially, academic research organizations. End-user representatives can sit on advisory boards, or make personnel available for consultation. They can allow researchers to have access to their facilities to interview workers and managers, observe work in progress, or conduct experiments in a realistic environment.

End-user organizations should also fund research into social applications, whether internally or in academia. Such new sources of funding are needed if the research agenda is to expand into social applications without detracting from continued research in the component technologies that are needed to keep progress in information technology going. Funding research will also allow end users to influence the research agenda and help produce strong returns on investment if the research is performed well. Many end-user organizations—from government mission agencies, such as the Internal Revenue Service and Federal Aviation Administration, to large corporations—make significant expenditures on the development of IT applications. Even modest improvements in their ability to design and deploy large-scale systems and social applications could result in significant development cost savings and more effective applications. Companies that operate large IT systems, including network operators and application service providers, will also see benefits from their investments in IT research, especially in large-scale systems research.

Organizations in particular application domains and vertical industries can also pool their resources to fund precompetitive research. They could also establish new research organizations patterned after some of the more successful existing examples like the Electric Power Research Institute, the Cable Research Laboratories, or the Semiconductor Research Corporation. This approach would allow better coordination of expenditures and create a critical mass of researchers addressing particular problems.

Perhaps the greatest obstacles to more interdisciplinary research and the greater involvement of end users in IT research are cultural. Today, top IT researchers are often unaware of the opportunities for making a major impact in application areas and of the serious and interesting challenges they face. Similarly, end-user organizations are often unaware of the opportunity for research to redress their serious challenges, or of the major opportunities arising from more effective use of IT that can be uncovered with research. Exceptions include the military, space exploration, and medicine, among other application domains, which long ago recognized the applicability of research (although not always related to IT) to their mission, and have realized tremendous benefits from these investments.

A hopeful precedent is scientific computing. There is a long history of successful collaboration between natural and social scientists and the IT–related disciplines (e.g., computer science and engineering, electrical engineering, and information sciences) centered on modeling, simulation, and data extraction and analysis. This collaboration has profoundly affected high-performance computing technology and is having a major impact on the practice of science itself. This collaboration was aided by a greater cultural affinity and mutual appreciation among computer scientists and other scientists.

Other hopeful historical precedents include mathematics, economics, and the social sciences, which have become increasingly prevalent components of many other disciplines. Like mathematics and economics, IT will increasingly become an integral part of many disciplines, with domain specialists acquiring a strong IT background. Conversely, the social sciences can become a more integral part of IT, just as social scientists have complemented engineering disciplines in areas such as urban planning, transportation, architecture, and structural engineering. All this is a natural—and hardly unprecedented—result of the increasing pervasiveness of IT.

Information technology is one of the most exciting and influential developments of our age, and together with other new technologies, it strongly underpins the success of the nation. In light of its increasing influence and impact, evident and growing problems in the application of IT must be taken seriously. The nation needs to invest increased resources for research, particularly in problem areas that are today largely ignored. The areas of large-scale systems and social applications both merit increased attention from researchers in application domains and in information technologies. Increased funding will be needed, especially from (or on behalf of) the end-user organizations and service providers that will be the most direct beneficiaries of research. The returns on investment could be enormous. As almost every citizen and organization depends increasingly on IT, the importance of improved understanding and better outcomes grows.

Acknowledgements

The Directorate for Computer and Information Science and Engineering of the National Science Foundation supported the NRC study summarized here. Other members of the study committee (the Committee on Information Technology Research in a Competitive World) contributed many of the ideas expressed in this article and in the full report, *Making IT Better* (NRC, 2000). Committee members Robert Sproull (NAE) and Stewart Personick (NAE) supplied valuable comments on this article.

Notes

1. The views expressed in this article reflect the personal perspectives of the authors and are not statements of the National Research Council, the National Academies, or the University of California, although they draw upon the authors' contributions to the NRC report, *Making IT Better.*

2. All economic data cited in this paragraph are derived from the U.S. Department of Commerce (2000).

3. The NRC's Computer Science and Telecommunications Board plans to release a study on IT workforce issues in the fall of 2000. Additional information is available online at http://www4.nationalacademies.org/cpsma/itwpublic2.nsf.

4. Estimates of industry and government spending on research and development related to information technology were made by members of the study committee that produced the report *Making IT Better*. The industry figures derive from data compiled for the National Science Foundation by the U.S. Bureau of the Census for the following six industries: office, computing, and accounting machines; communications equipment; electronic components; communications services; computing and data processing services; and wholesale trade of professional and commercial equipment and supplies.

References

Johnson, J. 1999. Turning chaos into success. Software Magazine 19(3):30–39. Available: http://www.softwaremag.com/archive/1999dec/success.html [11 August 2000].

Layton, L. 1999. Computer failure puzzles Metro: Opening delayed, rush hour slowed. Washington Post, 25 September, p. B1.

Leunig, E. 1999. Schwab suffers repeated online outages. CNET News.com (22 October). [Online.] Available: http://news.cnet.com/news/0-1007-200-922368.html [11 August 2000].

Meehan, M. 2000. Update: System outages top online brokerage execs' concerns. Computerworld (4 April). [Online.] Available: http://www.computerworld.com/cwi/story/0,1199,NAV47_STO43763,00.html [11 August 2000].

National Research Council (NRC). 2000. Making IT Better: Expanding Information Technology Research to Meet Society's Needs. Washington, D.C.: National Academy Press.

Standish Group International. 1995. The Chaos. [Online.] Available: http://standishgroup.com/visitor/chaos.htm [11 August 2000].

U.S. Department of Commerce. Economics and Statistics Administration. 2000. Digital Economy 2000. Washington, D.C.: Department of Commerce.

The Digital Dilemma

Randall Davis

The information infrastructure threatens the delicate balance between sharing information and protecting intellectual property.

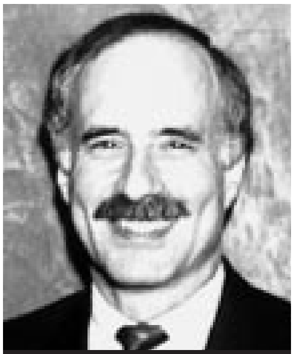

Randall Davis is a professor in the computer science department at MIT. He chaired the NRC committee that produced the report The Digital Dilemma, *upon which this article is based.*[1]

The combination of technologies that make up the information infrastructure—information in digital form, computer networks, and the World Wide Web—has arrived accompanied by contradictory powers and promises. For intellectual property (IP) in particular, the information infrastructure promises more—more quantity, quality, access, and markets—while simultaneously imperiling the rewards of those who create and publish. It is at once a remarkably powerful medium for publishing and distributing information and the world's largest reproduction facility, running unchecked in practice if not in statute. It is a technology that can enormously improve access to information, yet can inhibit access in ways that were never before practical.

That technology has arrived in a world where our existing laws, policies, and practices around IP depend on a number of subtle, surprisingly complex, and at times conflicting elements of law, public policy, economics, and technology. Those elements are in relative balance today, but may well be thrown out of balance by the transformations resulting from the information infrastructure.

One core of the problem is illustrated simply enough. A printed book clearly can be read by one or perhaps two people at once, people who must, of course, be in the same place as the book. But make that same text available in electronic form, and there is almost no technological limit to the number of people who can access it simultaneously, from almost anywhere on the planet. At first glance, this is wonderful news for the information consumer and for society—the electronic holdings of libraries (and friends) around the world can become available 24 hours a day, year round; they are never "checked out." These same advances in technology create new opportunities and new markets for publishers.

But there is also a more troublesome side. For publishers and authors, the question is, How many copies of the work will be sold (or licensed) if networks make possible planetwide access? Their nightmare is that the number is one. How many books (or movies, photographs, or musical pieces) will be created and published online if the entire market can be extinguished by the sale of the first electronic copy?

The nightmare of consumers, on the other hand, is that the attempt to preserve the marketplace leads to technical and legal protections that sharply reduce access to society's intellectual and cultural heritage, resources that have long been seen as crucial to democracy and science.

Balancing Control and Access

The information infrastructure thus has the potential to demolish the careful balancing of public good and private interest that has emerged from the evolution of U.S. intellectual property law during the past 200 years. The public good is the betterment of society that results from the constitutional mandate to promote the "progress of science and the useful arts"; the private interest is served by the time-limited monopoly (a copyright or patent) given to one who has made a contribution to that progress. The challenge is in striking and maintaining the balance, offering enough control to motivate authors, inventors, and publishers, but not so much control as to threaten important public policy goals, such as preserving the nation's cultural heritage, providing broad access to information, or promoting education, science, and scholarship. As usual, the devil is in the details, and by and large the past 200 years of intellectual property history have seen a successful, albeit evolving, balancing of those details. But the evolving information infrastructure presents a leap in technology that upsets the current balance, forcing a rethinking of many of our fundamental premises and practices.

The stakes involved in all this are high, both economically and in social terms. Decisions we make now will determine who will benefit from the technology and who will have access to what information on what terms—foundational elements of our future society.

Existing practices depend on complex elements of law, public policy, economics, and technology.

The difficulties we face today arise primarily from two sources: a trio of technological advances that produced the infrastructure and the emergence of computers as a routine part of everyday life. These two developments fundamentally altered the landscape, and their consequences present a significant array of challenges.

Of the many technological advances that led to the information infrastructure, three have particular significance for IP: the increasing use of digital information, the widespread reach of computer networks, and the creation of the World Wide Web.

Information in digital form is orders of magnitude easier, faster, and cheaper to reproduce than is information in analog form (e.g., hard copy). Digital copies are also perfect, so each copy can in turn become the seed for additional perfect copies, quite unlike the situation with traditional media such as photocopies.

These properties of digital information are for the most part widely acknowledged. Somewhat less well appreciated is the fact that accessing digital information inevitably means making a copy, even if only an ephemeral copy. This copying action is deeply rooted in the way computers work. For example, when you view a page from the World Wide Web, several copies are made—one sent from the remote computer to your computer, another copy stored on your hard drive, a third copy loaded into memory, and yet another copy displayed on the screen.

Such copying occurs with all digital information. Use your computer to read a book, look at a picture, watch a movie, or listen to a song, and you inevitably make one or more copies. Contrast this with the use of traditional media—reading a book does not involve making a copy of it, nor does watching a movie or listening to a song.

This intimate connection between access and copying has considerable significance in the context of intellectual property protection. One of the essential elements of copyright—the right to control reproduction—works as expected in the world of traditional media, where there is an obvious distinction between access and reproduction, and where the copyright owner's control of reproduction provides just that. But in the digital world, where no access is possible except by copying, complete control of copying would mean control of access as well, a consequence of considerable importance to all stakeholders.

Copying occurs with all digital information.

Placing information in digital form has consequences for distribution as well. Where information in hard-copy form is typically sold, digital information has traditionally been licensed. Packaged software, for example, has traditionally had a shrink-wrap license, an agreement by which the purchaser buys a right to use the software, but which provides no ownership in it. More recently, a wide variety of digital information is being licensed, including documents, databases, and images.

The difference between selling a work and licensing it is significant. The sale of a physical copy of a work has been the dominant model for transferring IP to the consumer for more than 200 years, and involves the complete transfer of ownership rights in a particular copy of the work. Copyright law explicitly anticipates the sale of intellectual property products and, by its "first-sale rule," gives the purchaser a significant body of rights in the purchased copy. The purchaser is, for example, free to lend, rent, or resell the purchased copy. In that sense, copyright law follows IP products into the marketplace and promotes the continued dissemination of information.

Licenses, by contrast, are contracts—private agreements between two parties that provide for limited transfer of rights to use an item. They can involve a wide range of terms and conditions, but, unlike copyright law, need not incorporate any public policy considerations beyond some basic limits on what constitutes an enforceable contract. To the extent that digital information is distributed by license, then, there is no statute, history, or tradition of incorporating public policy considerations such as fair use.

Computer Networks and the Web

Computer networks have radically changed the economics and logistics of information distribution, enabling information to be sent worldwide, almost for free, and (for items of reasonable size) almost instantaneously. The Web in turn has radically altered the economics and logistics of publication, allowing everyone to be a publisher with worldwide reach. Where reproduction and distribution put information in the hands of those who know they want it, publication makes people aware of information that is available, a function the Web performs well. The astonishing variety of documents, opinions, articles, and works of all sorts on the Web demonstrate that millions of people worldwide are making use of that capability.

This trio of technological developments—digital information, computer networks, and the Web—have together been the source of profound changes. Digital information radically changes the economics and character of reproduction, computer networks radically change the economics and character of distribution, and the Web radically changes the economics and character of publication.

For publishers, these technological developments have opened up new markets and new products, such as online music and books. But the same developments offer advantages to individuals or pirates making and distributing unauthorized copies: the process is orders of magnitude faster, easier, and less expensive. The important result is that the natural barriers to infringement have eroded significantly. Where unauthorized reproduction and distribution of hard-copy works is limited in part by the difficulty, expense, and loss in quality inherent in the process, with digital information, perfect copies can be made and distributed almost for free. Stakeholders on all sides of the issue wonder whether something can be put in place to restore the balance of forces, or whether the world has changed fundamentally and permanently.

The second major source of difficulties in the digital dilemma is the routine presence of computers and the Web in work settings, and increasingly in households as well. What was once in research laboratories is now a broadly available consumer product. One consequence is that individuals routinely have the means and opportunity to access and copy vast amounts of digital information, including software, text, and audio and video material, but have no clear picture of what is legal or acceptable. As a result, they are unprepared to deal with the IP issues the technology brings. Corporations dealing with these issues turn to their legal staffs, but individuals are bewildered, if indeed they are aware of the law at all.

A second consequence of the diffusion of the technology into everyday life is that intellectual property law and its enforcement are becoming increasingly concerned with private behavior. Copyright has traditionally been concerned with public actions with public consequences, such as public performance, public display, or the dissemination of copies, and it has focused on behaviors of organizations or individuals whose actions have large-scale public consequences. But with computer and communication equipment now commonplace in the home, individuals can do in private what once would have required substantial investment and perhaps criminal intent. As the potential impact of private behavior has grown, so correspondingly has interest in regulating that behavior. This shift in the focus of IP law represents an important consequence of information technology's emergence into everyday life and presents another social and policy challenge in managing the IP balance.

Increasing Use of Licensing

One of the problematic issues raised by the digital dilemma is the increasing use of licensing rather than sale as the primary mechanism for distributing information. Licensing is, as noted, familiar in the digital world in the form of software licenses, but is a newer phenomenon for most other digital information products. An increasing amount of the information acquired by libraries, for example, is in digital form, and unlike print materials, which historically have been available on a sale-of-copy basis, digital materials are frequently available only through licenses.

Licensing can have advantages—it may provide clarity on terms and conditions of access, and it may provide for increased rights for the institution that go beyond those provided under copyright (e.g., the ability to make unlimited copies for local use). Licensing may also increase the options for making information available. For example, a license may grant time-limited access to some part of a book or report, presumably for much less than the purchase price for the entire work.

But there are also concerns about licensing, particularly the impact it may have on public access. The increased use of licensing has the potential to diminish the public access accorded through the first-sale rule. Consider libraries as an archetypal example. In the print world, a library's failure to renew a subscription or buy an updated version of a book has no effect on the availability to patrons of earlier volumes or editions. In the world of licensed information, however, ending a subscription to an electronic journal may mean the end of access to earlier volumes or editions as well.

The natural barriers to copyright infringement have eroded significantly.

A second issue arises from the nature of licenses as contracts. Contracts need not incorporate, and indeed may attempt to override, the public policy considerations that have been carefully crafted into copyright law. Those who contract for digital information may find that their access is far more restrictive than what they were accustomed to for print materials, unless fair use and other such considerations are explicitly a part of the agreement.

Some institutions, notably libraries, have worked to negotiate licenses that preserve fair use and other public access features. Publishers are currently experimenting with licensing models to respond to these concerns. Yet the concern remains about the use of a mechanism such as licensing that lacks any of the built-in protections for public access embodied in copyright law.

Questions also arise about the interaction of licenses and copyright. For example, copyright law currently gives owners of copies of computer software the privilege to make backup copies. Can that privilege be taken away by a shrink-wrap license? Can a license term prohibit disclosing flaws to other potential users? These and related questions are far from resolved.

If licensing becomes the dominant means of distributing information in mass markets, additional concerns arise for works that are considered part of our intellectual and social heritage. One could imagine a world in which novels, poems, and paintings, for example, are available only (or mostly) by license rather than sale. The consequences of such a world for public access are far from clear.

Intellectual property law is becoming increasingly concerned with private behavior.

Finally, the trend toward licensing means that increasing amounts of information are delivered as experiences, rather than as artifacts. This is not entirely unfamiliar—where books are information artifacts, first-run movies are information experiences. We have lived with both for some time, but the difference matters. Buy a book and you own it forever; pay for access to a digital book and when the period of service is over you often retain nothing.

There are a variety of actions that can be taken in response to the challenges posed by the dilemma, a number of which are outlined in the National Research Council report, *The Digital Dilemma: Intellectual Property in the Information Age* (2000). We focus here on four: technological protection mechanisms, innovative business models, taking a broader perspective on the problem, and rethinking the nature of copyright.

The key technical problem in large-scale management of digital information is determining how to provide access without giving up all control. A variety of clever schemes have been proposed, many of which rely on a combination of encryption and rights management software. Encryption encodes information so that it can be accessed only with the appropriate key; rights management software enables fine-grained control of access, specifying such things as the number of accesses permitted, whether the material may be printed, etc.

A common scenario for using encryption involves making it both machine-specific and persistent. Encryption can be made machine-specific by incorporating into the encryption process some reliable property of the decryption/playback device. For a computer, for example, the serial number of the hard drive or CPU might be used. Then, even if both the file and encryption key are passed on, the information remains inaccessible because it cannot be decrypted on a machine with a different identifier.

But if the original purchaser, with authorized access, can decrypt the information, what prevents that person from passing on the decrypted file? Persistent encryption tries to address this, narrowing as much as possible the window of opportunity during which the decrypted information is available. In this approach information is decrypted just in time, that is, just before it is used, making it available as briefly as possible, and then only in small chunks at once. It is never stored, even temporarily.

Rights management involves providing some indication of the rights the consumer has purchased; this information is either encrypted along with the content itself or possibly maintained on a separate "rights server" accessed over the Web. Software on the user's machine then controls access to the information, consulting the rights listing to determine whether a requested action is permitted.

Limits of Technical Protection Methods

Technological solutions can be useful, but are limited in a number of important ways. One constraint arises from the need for consumer devices to be simple and fast. In some cases this precludes the use of industrial-strength encryption, which, though nearly impossible to crack, may be too slow for some consumer uses. The less powerful encryption systems used in commercial products, however, have routinely been cracked (e.g., the content scrambling systems used to encrypt DVDs).

Second, anchoring content to a specific machine raises an interesting problem for consumers—What happens when you upgrade your computer (or other playback device)? Must you repurchase everything you bought previously?

Finally, there is substantial difficulty in attempting to provide end-to-end protection within a general-purpose computer. PCs have been successful to a significant degree because they have open architectures; that is, components of the machine are accessible to the consumer (e.g., the hard drive, the video card, etc.). As long as the machine is designed this way, decrypted information can be intercepted and captured as it passes from one place to another inside the machine. Hardware and software designers could make such steps progressively more difficult, but the effort they would have to expend, and the consequential costs for all involved, would be substantial.

Innovative Business Models

In general, technology can play a useful role as a deterrent to unauthorized copying, but is far from a panacea. Consider, then, a second possibility: the use of innovative business models. By selecting an appropriate business model, a rights holder can at times significantly influence the pressure for and degree of illegal commercial copying and unauthorized reproduction by individuals.

Several general principles offer insight into the sorts of business models that can help. One principle suggests making the product cheaper and easier to buy than to steal. The basic point is to reduce the motivation to deal with unauthorized sources from the outset. Music provides one example: single tracks can be bought online for 99 cents (e.g., at www.emusic.com). Why pay a dollar for something you can get for free? Because, even with programs like Napster and Gnutella, the music isn't really free; there's an investment of time and effort in finding and in downloading it. With the informal, all-volunteer effort that is at the heart of all the file-sharing schemes, the difficulty of finding what you want may be substantial[2] and the time required to download (from an overloaded personal computer sitting in a dorm room somewhere) may be significant. Would you pay a dollar to avoid the hassle? Quite likely. Hence, even beyond the obvious issues of ethics and legality, the answer to the question "Why pay?" is, for service, reliability, and speed.

A second general principle suggests reconceptualizing the basic product. As noted, the digital music business may be primarily a service business, rather than a product business. As the value in the digitized music itself becomes increasingly difficult to protect (because digital information is so difficult to share without losing control of it), the value may reside in providing speed, reliability, and ease of access, rather than in artifacts like CDs or tapes.

The key technical problem is determining how to provide access without giving up control.

A third principle suggests finding the value in an alternative market. Again, music offers a plausible scenario. As testimony before Congress in July 2000 illustrated, the economics of the music business are not necessarily what one would expect, at least for the performers.[3] Some prominent and successful performers testified that during their entire careers they had received little or nothing in the way of royalties for their records, instead making the bulk of their earnings from live performances.[4] This suggests that, for some segments of the music industry at least, the music itself need not be the primary source of revenue; it might instead be a form of advertising for live performances.

These general principles suggest thinking about the digital information business in innovative ways and lead to a number of other nontraditional business models. Each of these models, importantly, creates an environment in which there is significantly less need for IP protection:

- Give away the product; make money from an auxiliary service. The Linux operating system, for example, is given away, yet a number of companies are in the business of providing Linux service, consulting, customization, and extensions.

- Give away the product; sell upgrades. Many antivirus vendors make available free, fully functional versions of the programs on their websites. They give away the program in order to sell subscriptions to the regular updates they make to the virus database.

- Give away one piece that promotes another. Adobe gives away their Acrobat Reader software to popularize the Acrobat PDF format and to create a market for all their other programs that create PDF files.
- Offer extreme customization. Custom CDs with a particular customer's selection of audio tracks are likely to be a less appealing target for reproduction.
- Offer a mass market product at a low price and high volume, with frequent improvements. Many software products fit in this category. Keeping the price low reduces the pressure for piracy, while constant improvements mean that the damage from unauthorized reproduction is time-limited.

Business models may offer a more effective means of dealing with IP issues.

None of these completely solve the problem, but each of them can sharply reduce the need for IP enforcement. There is also a more general point here about the relative power of law and business models: although legal prohibitions against copying are useful against large-scale pirates, they are unlikely to be nearly as effective against individual infringers, where detection and enforcement are problematic. Where such private behavior is concerned, business models may offer a far more effective means of dealing with IP issues.

The problems that arise from the interaction of IP and the information infrastructure need to be considered in a context that encompasses not only law, technology, and markets, but economics more generally, as well as psychology and public policy.

Importantly, each of these contributes its own approach and mindset. The legal view considers the law as it currently exists. Reference may be made to the legislative history and there may be appeals to analogies and past precedents, but the fundamental question is what the statutes and case law indicate as they now stand.

The technological perspective asks what mechanisms are available to address the problem (e.g., to inhibit copying) and considers questions of strength of protection, complexity of development, reliability, and so forth.

The economic perspective asks us to consider the possible role of markets, as well as the larger question of the costs and benefits of any proposed solution (including who pays the costs and who derives the benefits).

The psychological perspective concerns individual and group behavior—behavior that is grounded in perceptions of fairness, responsibility, fear, shame, guilt, convenience, and pragmatism, as well as in perceptions of the individual as beneficiary, victim, or patron in transactions involving intellectual property.

Finally, in the public policy view, the questions concern what goal society is trying to accomplish (e.g., is it better for society as a whole for such copying to be allowed or prohibited?), and might include a discussion of what the law ought to be in order to support important policy objectives (e.g., promoting national economic competitiveness).

This multiplicity of views is important in three ways. First, each of them brings a fundamentally different approach to the problem, phrasing the basic question differently. As all of them are relevant, there is power in considering each of these different conceptions of the problem. Second, some disagreements arise because the positions are grounded in different perspectives (e.g., law rather than economics), and are thus in effect asking (and answering) different questions. Finally, being aware of the multiplicity of perspectives may open up additional routes for dealing with issues; not every problem need be legislated into submission.

The Concept of Copying

One final response to the digital dilemma asks a fundamental question about intellectual property. All of the preceding discussion accepts a fundamental perspective that underlies copyright, namely, the concept of copying as a foundational legal and conceptual notion. As the very name of the law indicates, the right to control reproduction is central to copyright. Deciding whether a work has been copied has, as a result, been a fundamental question underlying much of copyright history and analysis.

But is the notion of copy still an appropriate conceptual framework in the age of digital information? Two reasons suggest it might not be. One reason, as we have seen, is that legitimate use of a digital work requires making a copy. Hence, noting that a copy has been made tells far less about the legitimacy of the behavior than it does in the hardcopy world, where there are very few legitimate rationales for copying an entire work.

A second reason is that, because copying is so bound up with the way computers work, controlling the act of copying, in the view of some, provides unexpectedly broad powers, considerably beyond those intended by the copyright law. In the world of physical works, once a work has been published, the rights holder cannot in any pragmatic sense control access to the copies distributed. Social institutions (such as bookstores and libraries) and individuals with copies enable any motivated reader to gain access to the information in the work.

But when access requires reproduction, the right to control reproduction is the right to control access, even the access to an individual copy already distributed. Authors would not, of course, routinely deny access to their published digital works. But because access requires reproduction, control of reproduction provides control of access to individual published copies, a right not conceived of as part of copyright and hence not to be embraced lightly, whether or not routinely exercised.

Two points are important here. First, in the digital world, reproduction loses much of its power as a predictor of important consequences, and hence the question of whether a protected work has been copied may be considerably less important. Second, control of reproduction in the digital world is a blunt instrument whose impact reaches considerably beyond the original intent, bringing into question its use in accomplishing the goals of intellectual property law.

Reconceptualizing Copyright

Considering that control of reproduction is a means, not the goal, can we find some other mechanism that is more tightly connected to the goal, whether in the digital or analog world? This will not easily be done, but one suggestion may help promote serious consideration of the issue.

It may be useful to start not by asking whether a copy has been made, but from considering what the law is attempting to achieve—ensuring progress in the sciences and arts—and ask instead whether a use being made of a work is substantially destructive of a common means of achieving that goal, namely, providing incentive to authors. This approach is similar in overall spirit to the concept of fair use, which requires consideration of the impact on the market for the work or on the value of the work. But it is somewhat broader in scope, as incentive arises for authors in more than the marketplace alone, coming as well, for example, from the ability to control the time, place, and manner of publication.

This view would not conflict with all of the other traditional exclusive rights in copyright. Creation of derivative works, distribution, public performance, and display of the work can all be conceived of and protected on grounds independent of whether a copy has been made. They also have impact on incentive, whether via economic effects in the marketplace or other factors, and hence would be consistent with an incentive-based analysis.

Is the notion of copy still an appropriate conceptual framework?

Any such substantial change would of course also bring problems. There would have to be a substantial period of familiarization and a means of dealing with the tension between trying to make such a law easier to follow (by drawing a sharp line defining what constitutes incentive-destroying use) and keeping the criteria more general, as is the case for fair use, so the law can address unanticipated situations in the future. Nevertheless, the discussion and the examination of what constitutes a well-grounded model of IP protection in the digital world may well be worth the effort.

The development and deployment of the information infrastructure presents a variety of challenges, ranging from the pragmatics of enforcing laws that can be casually broken by individuals (in private, inexpensively, and almost undetectably) to the possibility and perhaps the need to rethink some of the foundational concepts underlying those laws. Yet, from its origins in patent laws of Venice in the 1400s through to modern times, intellectual property law has evolved and changed in response to such challenges. What is likely to happen this time? Several things.

To some degree, society will simply adjust to the new reality and, in part, carry on in familiar ways. Recall how software vendors gave up on the awkward technical mechanisms used in the 1980s to defeat piracy (e.g., distribution disks that could not be copied in the ordinary way), and continued to do business and prosper in a world where there is nontrivial piracy. So it is likely to be with the distribution of digital content. Customers will grow used to subscriptions to online information, and authors and publishers will continue to do their work in the presence of some unauthorized reproduction.

But as the turmoil in the music, publishing, and movie industries suggests, the upheaval created by new technology may mean that accommodation is not enough. New business models will need to be explored and tested. New approaches to IP issues will need to be founded on more than law and technology, embracing as well an understanding of the economics of information, sociology, and psychology. And IP itself may need to be conceptualized to some degree, in recognition of the changes we now face.

The issues are difficult, but inescapable. We are all unavoidably engaged in an experiment that tests our ingenuity and resourcefulness in finding ways to accomplish the apparently paradoxical goals of motivating individual creation while still reserving the ultimate benefits of that creation for the good of society, in a world where replication, distribution, and publication are astonishingly easier and less expensive than they have ever been. The issues are difficult, and important—the decisions we make will shape foundational elements of our future society.

Notes

1. This article draws heavily on the text of *The Digital Dilemma: Intellectual Property in the Information Age* (2000), a report produced by the National Research Council's Computer Science and Telecommunications Board. The author chaired the study group, the Committee on Intellectual Property Rights and the Emerging Information Infrastructure. The study was funded by the National Science Foundation.

2. Looking for a song by Neil Young and can't find it? Try looking up "Niel Young."

3. U.S. Senate Committee on the Judiciary, "The Future of Digital Music: Is There an Upside to Downloading?," 106th Congress, 2nd session, 11 July 2000.

4. See, for example, "Byrds Man Sees Promise of Digital Music," *Boston Globe,* 14 July 2000, p. D1.

References

National Research Council. 2000. The Digital Dilemma: Intellectual Property in the Information Age. Washington, D.C.: National Academy Press.

Wiretapping the Net

Steven M. Bellovin

To serve the needs of law enforcement while protecting privacy, the legal and technical approaches to Internet wiretapping must be reexamined.

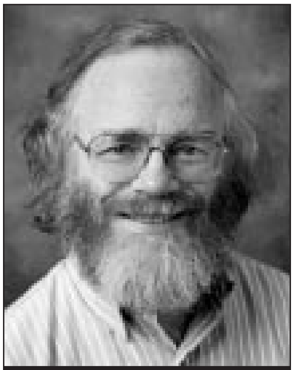

Steven M. Bellovin is a fellow at AT&T Labs Research, where he does research in networks and security.[1]

Of late, there has been a great deal of interest in the ability of law enforcement agencies to place wiretaps on the Internet. While there are certainly legitimate reasons for wanting this ability, it is an area that is fraught with technical difficulty and legal ambiguity. In light of these problems, new legal and technical approaches to wiretapping are in order.

In this article, we examine wiretapping problems from several perspectives—statutory, jurisdictional, and technical—and suggest paths to minimizing those problems.

First, and in some sense the simplest, are problems with current statutes. While some of these can be fixed by legislative action, others raise deeper issues. Second are jurisdictional problems; Internet routing systems can make it unclear who has the right to tap a call and under what circumstances. Third, the very nature of Internet communications introduces complex technical problems. Packet-switched networks are inherently much harder to monitor than conventional phone lines. When looked at in combination, these factors highlight an overall complexity that makes the practice of Internet wiretapping a dubious undertaking.

Before discussing legal issues, it is important to understand, on a basic level, how the Internet works. The Internet is a packet-switching network; that is, a stream of data sent from one computer to another is split up into small pieces known as packets, and each packet is transmitted independently. Because each packet travels independently, all packets must be labeled with their source and destination addresses. These are known as IP (Internet Protocol) addresses and are not normally seen by end users.

New legal and technical approaches to wiretapping are in order.

Packets are transported over the Internet by a series of hops between devices known as routers, which serve as intermediate network nodes. A router looks at a packet's destination address and forwards it to the appropriate neighboring router. The most appropriate neighbor is determined dynamically by the routers—they regularly exchange information about the network topology by means of routing protocols. If two or more neighbors are equally good, a router will choose a neighbor so as to balance the load between alternate paths.

Each ISP (Internet service provider) operates its own network of routers. ISPs talk to each other via private links and at public exchange points where many ISPs meet. These interconnections, especially the public exchange points, are often overloaded, and the resulting congestion is a major source of data transmission errors, known as dropped packets.

There are no guarantees of packet delivery on the Internet—packets may be dropped, duplicated, damaged, or reordered along the way. All users who require reliable packet delivery must make arrangements for error detection and correction. Of course, most do, usually by using TCP (Transmission Control Protocol) in conjunction with IP. TCP detects damaged or dropped packets and sends acknowledgements back to the sending system when good packets are received. If the sending system does not receive the proper acknowledgements, it will retransmit the missing pieces until they arrive intact.

Applications that do their own retransmissions or that do not need all of the functionality of TCP sometimes use UDP (User Datagram Protocol). UDP is considerably lower in overhead and is often used for simple query/response applications.

Internet Addressing

In general, ISPs assign IP addresses to their clients on a dynamic basis. This is partly because addresses are in short supply and partly because IP addresses need to be flexible to reflect the current topology of the network. Consumers who use dial-up modems can end up connecting to the Internet through different routers (sometimes in other cities), depending on the load on their ISP's local modem pool.

Internet servers tend to have reasonably constant IP addresses, and they always have well-known names. These names are mapped into IP addresses via the DNS (Domain Name System). Apart from Web servers, ISPs run a number of servers on behalf of their users, notably e-mail and netnews machines. These servers are often replicated to provide load-sharing and reliability, and the duplicates are often geographically remote from the primaries. Similarly, corporate servers are often connected to more than one ISP, each of which could assign its own IP addresses to the corporate servers.

Internet connections are between pairs of systems, but performing user-requested services may involve intermediate systems. Consider, for example, e-mail sent between two typical home users. The mail is initially sent from the first home computer to the ISP's outgoing mail gateway. From there the mail is sent to the receiving ISP's incoming mail gateway, and from there it may be forwarded to a mail repository server. Finally, the receiving user dials in to the ISP, connects to the mail repository, and downloads the mail. At least three, and probably four or more, separate TCP connections are involved, as well as several DNS lookups. The multiplicity of systems involved in carrying out even simple requests is at the root of some of the legal complexities.

Any legal wiretapping has to be done in accordance with statutory authority. While in some sense this is an easy problem—the appropriate legislative bodies can simply enact any necessary laws—the problems of definition are significant. In particular, concepts familiar to law enforcement in the traditional telephony world do not necessarily translate easily to the Internet.

Statutory Considerations

The basic framework of U.S. wiretap law was adopted in 1968 as Title III of the Omnibus Safe Streets and Crime Control Act. (As a consequence, law enforcement personnel often refer to court orders permitting wiretaps as "Title III orders.") The law was significantly amended in 1986 by the Electronic Communications Privacy Act (ECPA), the primary thrust of which was to add the anti-eavesdropping protections—and the wiretap permissions—from the voice world to the data world.

As the law stands now, wiretap permissions are governed by four major sections of statute: 18 USC 3121 (pen registers and trap-and-trace devices); 18 USC 2510 (interception of communications); 18 USC 2701 (access to stored communications); and 50 USC 1801 (foreign intelligence surveillance).

The pen register statute poses the greatest statutory problem. Pen registers are devices that record what telephone numbers are dialed *from* monitored lines, while trap-and-trace devices record the phone numbers that dial *to* a monitored line. These concepts were once well defined, but they pose considerable problems when extended to the Internet.

For example, what are the Internet equivalents to dialing and dialed numbers? The most obvious answer is the IP addresses in each packet, which represent the actual endpoints of the communication. But these endpoints are likely to be either uninteresting or well beyond the scope of a reasonable warrant. With a typical e-mail server system, all that could be learned by monitoring an end user's line is that the user is sending or receiving e-mail; the identity of the user is not disclosed. Similarly, putting the monitor at the ISP would reveal that a customer of one ISP was corresponding with a customer of another ISP. Neither party's identity would be revealed by monitoring at this level.

A more useful answer, from the perspective of the information to be learned, would be to monitor the actual e-mail addresses used. This, too, is problematic, for a number of reasons. First, e-mail addresses are not authenticated. It's quite easy to supply false source addresses by simply changing the mailer's configuration. This is a strong argument against trap-and-trace monitoring based on the sender's address. Similarly, destination addresses can be fabricated, and in this case, delivery failure notification would be e-mailed to the sender, and may not be detected by the monitor.

A more compelling problem is that the necessary user information may not exist, or it may be in the wrong place. E-mail sent via a "bcc:" option does not contain the recipient's name in the mail header lines, and typical mail retrieval protocols do not distinguish between mail header lines and mail content. Indeed, that lack of differentiation raises the most serious technical issue.

Internet standards have distinguished between the "envelope" and the "contents" of e-mail since at least 1982. In fact, those precise words are used. The envelope contains the instructions to the mail system about the sender and the recipients. With a pen register warrant, is it legal to look beyond the envelope? And if it were legal to look, the content in header lines poses two further problems. First, the information in a header does not necessarily relate to the actual sender or recipients of the mail. It is quite easy to manipulate a header to list addresses for people who will never see the mail, leading to the possibility of innocent parties being dragged into an investigation. Second, header lines may reveal important information about parties not covered by the court order—the addresses of other correspondents of the sender—thus going well beyond the capabilities of a traditional telephone trap-and-trace device. That is, if someone sends e-mail to the target of an Internet trap-and-trace order, the sender's identity will be disclosed to the investigators, as intended by the court order. But other addresses listed in the headers will be disclosed as well, despite the lack of any statute intending this result.

What are the Internet equivalents to dialing and dialed numbers?

A final problem to solve is the wording of the statute itself. Currently, a pen register is defined as "a device which records or decodes electronic or other impulses which identify the numbers dialed or otherwise transmitted on the telephone line to which such device is attached.” Similarly, a trap-and-trace device is defined as “a device which captures the incoming electronic or other impulses which identify the originating number of an instrument or device from which a wire or electronic communication was transmitted.”

Many of the technical problems with wiretapping stem from the very nature of Internet technology.

Note that both definitions not only specify a telephone line, they specify a “device” or a “line.” These concepts do not correspond to an identity, or even an e-mail account. At best, the analogs are a computer and either the access line or the IP address assigned to the computer. As noted, however, the latter is subject to dynamic change and is not particularly useful when dealing with e-mail.

The legal authority for full-content wiretap warrants is more clear. The ECPA amended the statute to speak of “electronic communications,” rather than just voice calls. Furthermore, the distinctions between IP address and destination, or between message envelope and content, are irrelevant; the investigator is entitled to all traffic. The difficulty comes in identifying the content belonging to the targeted user.

Most physical media used to carry Internet traffic are shared, with the exception (in some cases) of the access line to a customer’s premises. That is, the same physical wire or fiber-optic cable carries traffic to or from many different parties. To isolate a particular party’s packets, it’s necessary to look at the IP addresses. Is this sort of examination legal?

This question leads to questions of jurisdiction—Who has the right to place a wiretap, and under what circumstances? Does the physical presence of a packet in some particular locale matter?

Packets on the Internet can take a complex path from source to destination. This is partly due to the nature of IP routing, but even more to the complex relationships among ISPs. It is rather rare for a conversation to stay within a single ISP; just how and where they interconnect is governed by complex business and technical considerations.

A few recent experiments by the author make clear just how nonintuitive routing can be. In one case, packets between two towns in North Carolina went via Atlanta, Georgia. Packets from the author’s office to his home, both in New Jersey, went via California. And packets from New Jersey to Russia went to New York, Washington, D.C., and back to New York before finally heading overseas.

Access to e-mail raises more troubling issues. As noted, most e-mail addressed to individuals will reside on an ISP’s data center, on a mail server, until explicitly retrieved by the recipient. It is not likely that the data center will be in the same jurisdiction as either the sender or the recipient. What judge has the power to order access to such messages?

This problem is further complicated for international traffic. For example, in one test, traffic from North Carolina to Costa Rica went via Montreal. Does that give Canadian authorities any right to read it? For historical reasons, the United States is in the middle of many Internet paths. Does this give the United States the right to read such traffic?

Unreliable Packet Switching

Many of the technical problems with wiretapping stem from the very nature of Internet technology. Someone who wishes to avoid monitoring can exploit the complexity of the technology.

For example, one problem inherent to the Internet is that of packet stream reassembly. The individual packets that make up a message must be recombined at destination to form a coherent whole, and the rules for doing this are complicated. If the process is implemented differently on the monitoring box and the recipient’s system, the two might see different streams, especially if the target user attempts to evade the monitor. For example, consider two packets whose contents overlap in the final stream. This is acceptable to TCP, and in some cases is a normal occurrence; TCP compares the overlapping areas and ignores the duplicate content. But what if the two areas differ? Which packet should the monitor believe?

An attacker can make the problem even worse by exploiting packet lifetimes. Packets have a finite lifetime, measured in router hops. The "hop count" is assigned by the sending system, and each router on the path subtracts one hop. If the count reaches zero, the packet is discarded. Suppose there is a sequence of packets containing a login name, a set of backspace characters, and a different login name. Which login name is intended for the recipient? If the packets all have the same lifetime, the second one would be used. But if the backspace characters and second login name have too short a lifetime, the first name would be used. Can the monitoring system handle it properly? Using purely passive techniques, it is very difficult to tell how far away a destination is, and simply seeing different lifetimes on different packets says nothing about whether or not a destination will receive them.

Packets to a given destination can take different paths through the Internet. This can reflect topology changes or load balancing. Indeed, even a single computer can use multiple dial-up sessions in parallel to achieve greater throughput. A simple monitoring station may not be equipped to detect this. Furthermore, a very high percentage of paths are asymmetric: return traffic does not flow through the same routers as forward traffic.

Even deciding which packets to monitor is difficult. As noted, consumer machines generally have dynamic IP addresses. A monitoring station needs to know what IP addresses to watch, and that means it has to monitor the address assignment protocol, which can be difficult to do. If the monitoring station misses assignment messages, it will not begin to monitor the target; if it misses disconnect messages, it will record someone else's traffic as well.

In fact, it is sometimes impossible to know what address is being used by a target. For example, some systems use Network Address Translators (NATs) to dynamically map a group of internal, private addresses to a few external IP addresses. Because of the shortage of IP addresses, some ISPs and many hotels employ NATs. A box monitoring a system on the public Internet has no way of learning the actual IP address of a correspondent system behind an NAT.

Software Complexity

The preceding description makes it clear that any monitoring system will of necessity be quite complex. That complexity carries with it its own risks—the most important being that complex software is buggy. It is generally accepted that the number of bugs in a system increases roughly as the square of the size of the code, and while bugs are never good, their consequences can be especially serious in an Internet wiretapping device.

The most obvious risk, of course, is that the device will crash. In some sense this is not so bad, in that such a failure is relatively obvious and benign, although it still represents wasted resources. More importantly, reliance on an Internet wiretapping device can divert investigators from the use of other techniques, and if the wiretap fails, no information will be gathered, by any means.

To make Internet wiretaps more accurate and secure, it is critical to limit vulnerabilities due to software bugs.

More subtle failures can have more serious consequences. Failures to record certain classes of traffic can easily deceive investigators; both exculpatory and incriminating evidence can be missed. Corruption in recording is worse yet. Apart from anything else, a recording that is demonstrably inaccurate is useless at trial, especially if it contains extraneous traffic.

But by far the most serious failure mode would be a takeover of the monitoring box by hostile parties, a scenario that is not at all improbable. About half of all new security failures are caused by "buffer overflows." If a buffer overflow were to be found and exploited in an Internet wiretapping device, the device itself could be taken over, possibly by the target of the wiretap. And the consequences of that—the potential for criminal control of law enforcement tools—are chilling.

If nothing else, the attacker would be able to learn the monitoring parameters, and, depending on the design of the monitoring box, might be able to alter or erase logs of previously recorded sessions. Worst of all, the device—a system that by design is a high-quality wiretapping unit—could be diverted and used as an eavesdropping unit for the attacker.

We have outlined a number of difficulties involved with wiretapping on the Internet. While some of the problems are very hard to solve, we can address some issues in a number of respects.

The first approach, of course, is to clarify the statutes. While the ECPA was a good first try, experience has shown that it does not match the reality of the Internet. The problem of pen registers, in particular, is a thorny one, given the inherent difficulty in determining the endpoints of the conversation—the target and the target's correspondent—in a way that cleanly separates that information from the content of their communications.

A second approach is to push the wiretap as close as possible to the target. Much of the trouble arises from differences between what the user sees and what the monitoring box sees. Other problems come from identifying just which packets belong to the user. A modem tap on the physical phone line would finesse many of these issues; packets on that wire are, by definition, to or from the user. Questions of which IP address to monitor are moot. Solutions of this nature (placing the tap near the target) could also be applied to DSL (digital subscriber line) connections, but not to cable modems, which are inherently shared.

Finally, simplifying the task to be performed will simplify the software. (That is, if the tap is closer to the target, there are fewer variables to handle with software. Among other things, it's much easier to monitor a slow line than a fast, multiuser cable.) Software complexity is the greatest unsolved problem in the computer industry and is likely to remain so. To make Internet wiretaps more accurate and secure, it is critical to limit vulnerabilities due to software bugs. Only in this way can we be confident that both the restrictions and the authorizations of the law are carried out.

Notes

1. The author recently served on a committee of the National Research Council that produced the report *Trust in Cyberspace* (National Academy Press, 1999). The Committee on Information Systems Trustworthiness operated under the NRC's Computer Science and Telecommunications Board.

NAE News and Notes

NAE Newsmakers

Leo Baranak, president of BBN Technologies; **George Heilmeier,** chairman emeritus of Telcordia Technologies; and **John Whinnery,** professor emeritus at the University of California, Berkeley, were named as **Eta Kappa Nu Eminent Members** during the IEEE Honors Ceremony on 24 June in Vancouver, B.C. They were recognized for technical contributions that have significantly benefited humankind.

Arthur E. Bergles, Clark and Crossan Professor of Engineering, Emeritus, Rensselaer Polytechnic Institute, received the **Luikov Medal** of the International Centre for Heat and Mass Transfer and honorary degrees from the University of Oporto, Portugal, and Rand Afrikaans University, Johannesburg, South Africa. He has also been appointed Glenn L. Martin Institute Professor of Engineering at the University of Maryland, College Park.

Y. Austin Chang, Wisconsin Distinguished Professor, University of Wisconsin, received the **John Bardeen Award** in March 2000 during the annual meeting of the Minerals, Metals & Materials Society. Dr. Chang was recognized for his seminal contributions to understanding of metal/compound semiconductor interactions.

James M. Coleman, Boyd Professor, Louisiana State University and Agricultural and Mechanical College, was elected a member of the **Russian Academy of Natural Sciences.** Dr. Coleman also received the **Kapitsa Medal of Honor** for his contributions to the field of petroleum sciences.

Albert A. Dorman, founding chairman, AECOM Technology Corp., received the first **Outstanding Lifetime Achievement in Leadership Award** from the American Society of Civil Engineers on 29 April. He was recognized "for his exception leadership skills in civil engineering demonstrated through his more than 50-year career, which has been distinguished by the size and diversity of the engineering and architectural operations he has managed."

Robert A. Duffy, president and director emeritus, The Charles Stark Draper Laboratory, was inducted as a **Fellow of the Institute of Navigation** for his contributions to the field of guidance, control, and navigation.

John A. Focht, Jr., senior consultant, Focht Consultants, Inc., was elevated to **National Honor Member** of the Chi Epsilon Civil Engineering Honor Society on 4 March 2000. He was honored as an authority on foundations for high-rise buildings in deep soil areas, foundations for offshore structures, pile foundations, and earth dams.

Alexander F. Giacco, CEO of Rheometric Scientific; **Christopher C. Kraft, Jr.,** independent consultant from Houston, Texas; and **James E. Turner, Jr.,** former president and CEO of General Dynamics Corp., were inducted as members of the **Academy of Engineering Excellence** of Virginia Tech's College of Engineering and its Advisory Board. This honor is reserved for individuals holding an engineering degree from Virginia Tech who have made sustained and meritorious engineering and/or leadership contributions during their career.

Solomon W. Golomb, Andrew and Erna Viterbi Professor of Communications and University Professor, University of California at Los Angeles, received the **IEEE Richard W. Hamming Medal** "for fundamental contributions to the theory of shift register sequences and their applications in digital communications."

Andrew S. Grove, chairman of Intel Corp., was given the IEEE's highest award, the **IEEE Medal of Honor.** Dr. Grove is being recognized "for pioneering research in characterizing and modeling metal oxide semiconductor devices and technology, and leadership in the development of the modern semiconductor industry."

Martin C. Hemsworth, retired senior consultant, GE Aircraft Engines, received the ASME International's **R. Tom Sawyer Award** in May for contributions to gas turbine technology spanning 50 years, including the revolutionary TF39 high-bypass engine used for the C5A large military transport and the derivative CF6 engines that power the vast majority of commercial aircraft worldwide.

John L. Hennessy, provost, Stanford University, and

David A. Patterson, E. H. and M. E. Pardee Professor of Computer Science, University of California, Berkeley, received the **IEEE John Von Neumann Medal** "for creating a revolution in computer architecture through their exploration, popularization, and commercialization of architectural innovations." Dr. Patterson also received the **IEEE James H. Mulligan, Jr., Education Medal** "for inspirational teaching through the development of creative curricula and teaching methodology, for important textbooks, and for effective integration of education and research missions."

Yu-Chi Ho, Gordon McKay Professor of Engineering and T. Jefferson Coolidge Professor of Applied Mathematics, Harvard University, was recently elected a foreign associate of the **Chinese Academy of Sciences** and elected to the **Chinese Academy of Engineering.** Dr. Ho also received the 1999 American Automatic Control Council's **Bellman Control Heritage Award** and the 1999 American Society of Mechanical Engineers **Oldenburger Award.**

John P. Holdren, Teresa and John Heinz Professor of Environmental Policy and director of the Program in Science, Technology, and Public Policy, Harvard University, was awarded the 2000 **Tyler Prize.** Dr. Holdren was honored for the significant role he has played in mobilizing the international community of scientists and policymakers to take action on a wide range of global energy, environmental, and security issues.

Amos E. Joel, Jr., was awarded the **IEEE Eminent Members' Award** for his developments, teaching, and writing in the field of electronic switching systems.

Ralph Landau, consulting professor of economics and senior fellow, Stanford University, received the 2000 **Petrochemical Heritage Award** at the International Petrochemical Conference in San Antonio, Texas. Dr. Landau was cited for his contributions and leadership in development, engineering, and marketing of new chemical processes.

William D. Manly, consultant, Oak Ridge National Laboratory, received an **Honorary Doctor of Engineering** from Notre Dame University on 21 May 2000. He was honored for being an international leader in the development of advanced high-temperature materials and associated processing technologies and because he has been at the forefront of technology transfer and the cross-fertilization of technology in American industries.

Arthur A. Oliner, professor emeritus, Polytechnic University, Brooklyn, N.Y., was chosen to receive the **IEEE Heinrich Hertz Medal** "for many outstanding contributions to the theory of guided waves and antennas, with emphasis on the fundamentals and applications of leaky waves."

R. Wayne Skaggs, William Neal Reynolds Professor and Distinguished University Professor, North Carolina State University, Raleigh, has been elected **president of the American Society of Agricultural Engineers** for the year 2001–2002.

Merrill Skolnik, superintendent emeritus, Radar Division, Naval Research Laboratory, has received the **IEEE Dennis J. Picard Medal for Radar Technologies and Applications.** Dr. Skolnik got the award "for outstanding leadership of Navy radar research, authorship of widely used books on radar, and personal contributions to the advancement of radar and of radar technology and systems."

John B. Slaughter, Irving R. Melbo Professor of Leadership in Education, University of Southern California, has been named the new **president and CEO** of the National Action Council for Minorities in Engineering.

Ponisseril Somasundaran, director, NSF/IUCR Center for Surfactants and La Von Duddleson Krumb Professor, Columbia University, was honored by a special symposium on Surfactants and Polymers at Interfaces during the 13th International Symposium on Surfactants in Solution held in June. Dr. Somasundaran was recognized for his landmark contributions in the area of surface and colloid chemistry.

At graduation exercises on 8 May, the University of Nebraska, Lincoln, presented the **Distinguished Service Award** to **William E. Splinter,** who retired as vice chancellor for research in 1993. He currently serves as volunteer director of the L. F. Larsen Tractor Test and Power Museum.

Dean E. Stephan, retired president, Charles Pankow Builders, was selected as the winner of the 2000 **Henry L. Michel Award for Industry Advancement of Research** by the Civil Engineering Research Foundation (CERF). The Michel Award recognizes leaders of the design and construction industry. Mr. Stephan is "recognized as a leader in the application of advanced concrete technology" and received the award on August 15 at the CERF Global Innovation Dinner in Washington, D.C.

Chang-Lin Tien, University Professor and NEC Distinguished Professor of Engineering, University of California, Berkeley, has been named the winner of Tau Beta Pi's **Distinguished Alumnus Award.** Dr. Tien is honored for his numerous and far-reaching contributions to the betterment of society, his passionate commitment to academic excellence, his role in improving international and intercultural relations, and his promotion of the standards of excellence in the engineering profession.

• • •

The following NAE members were **elected members to the National Academy of Sciences (NAS)** during the 2000 NAS Annual Meeting: **Howard Brenner,** Willard H. Dow Professor, Massachusetts Institute of Technology; **Robert L. Byre,** professor of applied physics and director, Center for Nonlinear Optical Materials, Stanford University; **Jean M. J. Frechet,** professor, University of California, Berkeley; **Thomas Kailath,** Hitachi America Professor of Engineering, Stanford University; and NAE foreign associate **Roddam Narasimha,** director, National Institute of Advanced Studies, Indian Institute of Science.

NAE Elects Officers

The National Academy of Engineering (NAE) has elected a new chair, home secretary, and three councillors as members of the Academy's governing council.

George M. C. Fisher, chairman of the board, Eastman Kodak Co., Rochester, N.Y., will serve a two-year term as chair of the NAE Council. Before joining Eastman Kodak, Fisher was chairman and chief executive officer of Motorola Inc., and spent 10 years in research and development at Bell Laboratories. The chair works with the president of the NAE to publicly promote the Academy and its policies. Fisher was elected to the NAE in 1994.

W. Dale Compton, Lillian M. Gilbreth Distinguished Professor of Industrial Engineering, and interim head of the School of Industrial Engineering, Purdue University, West Lafayette, Ind., will serve a four-year term as the Academy's home secretary. Before joining the Purdue faculty, Compton spent two years at the NAE as its first senior fellow (1986–1988); prior to that he was vice president of research at Ford Motor Co. Compton, who was elected to the NAE in 1981, will oversee the Academy's membership activities.

Newly elected as councillor for a three-year term are **Ruth M. Davis,** president and chief executive officer, Pymatuning Group Inc., Alexandria, Va., and chairman of the board of the Aerospace Corp., El Segundo, Calif., and **Paul E. Torgersen,** John W. Hancock Chair of Engineering and former president, Virginia Polytechnic Institute and State University, Blacksburg. Reelected to a second three-year term as councillor is **Delon Hampton,** chairman and chief executive officer, Delon Hampton & Associates, Washington, D.C., and president of the American Society of Civil Engineers.

In accordance with the Academy bylaws, a fourth councillor was chosen by vote of the Council at its May meeting, to ensure that the distribution of engineering disciplines on the Council are representative of the Academy membership. The fourth councillor is **Michael P. Ramage,** executive vice president, ExxonMobil Research and Engineering Co. The terms began 1 July 2000.

Also effective 1 July are the terms of six new NAE section chairs. The new chairs include **John H. Seinfeld,** section 3, chemical engineering; **John G. Bollinger,** section 8, industrial, manufacturing and operational systems engineering; **James C. Williams,** section 9, materials engineering; **Richard J. Goldstein,** section 10, mechanical engineering; **Robert R. Beebe,** section 11, petroleum, mining and geological engineering; and **Essex E. Finney, Jr.,** section 12, special fields and interdisciplinary engineering. The other six section chairs, with continuing terms, include **Steven D. Dorfman,** section 1, aerospace engineering; **Van. C. Mow,** section 2, bioengineering; **Loring A. Wyllie, Jr.,** section 4, civil engineering; **Robert E. Kahn,** section 5, computer science and engineering; **John F. Ahearne,** section 6, electric power/energy systems engineering; and **Frederick H. Dill,** section 7, electronics engineering.

From the Home Secretary

Simon Ostrach

As I start to write this, my last column as home secretary, I must reflect on the many years of my association with the NAE. I was elected to the Academy in 1978, as one of 100 members elected that year, when the size of the Academy totaled only 857 members and 58 foreign associates. During the ensuing years I served on almost every NAE committee and was chair of each of them. These experiences not only enabled me to become familiar with the operations of the Academy but also gave me an appreciation of the professionalism and dedication of the staff.

In 1992, when I was elected as the third home secretary in the Academy's history, succeeding Al Flax, the membership had grown to 1,626 members and 131 foreign associates. I was somewhat apprehensive in taking on the responsibilities, because how does one replace a person as remarkable as Al Flax? Al was most helpful in making the transition and even provided me with a notebook of "typical" responses to members' correspondence. I quickly learned, however, that there never was much that was "typical" in interacting with members. This led me to the conclusion that on any issue the members' responses followed a Gaussian distribution, and each element was certain it was correct. This theorem was validated throughout the years and it made the job most interesting, if daunting. My primary goals were to respond to each and every member's concern directly in a polite and serious manner, to make the election process as fair as possible and representative of all branches and types of engineering, and to try to obtain a reasonable distribution of members from all geographical regions and from underrepresented groups.

During the four years of my first term, some major transitions occurred in the Academy that called for opening the communication channels to the members and allowing them to participate more directly in the Academy's operations, including the election process. An NAE section liaison system was implemented to interact with the NRC communication structure as a first step in that direction. Those challenging times enabled me to set a precedent which, hopefully, will not be repeated, namely, of the home secretary working with three NAE presidents.

The major thrust of my second term was to have the first major review in more than 15 years of the membership structure, election procedures, and Academy size. A membership task group was formed to implement the review, and their recommendations, after considerable input from the membership, were adopted. These are directed to be responsive to many of the members' concerns and to engage the members more directly in the election process by such means as closer coupling of the peer committees and the sections, electronic versions of forms used in election cycles, and more information on the NAE website.

During the eight years of my terms as home secretary, the Academy has continued to grow and develop. Today, after the February 2000 election, there are 2,027 members and 157 foreign associates. In many ways the Academy is very different from what it was when I was elected and even from when I first took office. The policies and procedures are better known to the members, and the members play a greater role in all the operations. This is evident from the significant reduction in the correspondence containing complaints or asking for clarification of many issues. I am particularly pleased that 29 of the current 55 women members were elected during my tenure.

As an NAE officer and member of the NRC Governing Board and its Executive Committee, I had the opportunity to interact with and have close contact with the officers and Council members of the NAS and the IOM, as well as many of the NRC staff. All of these activities enabled me to meet, interact, and befriend many different people from all walks of life, people I would never have had the opportunity to know otherwise.

My experiences were rich and rewarding and I am grateful for the opportunity to have served such a unique and prestigious institution. There is no other organization in which there is such a diversity of talented people from every sector of the workforce who interact so freely. This is the strength and distinctiveness of the NAE and the basis upon which we can continue to grow and develop in our service to the nation and to our members.

It has been a great ride and whatever I may have contributed could not have been done without the support and cooperation of the officers, councillors, many members, and people throughout the National Academies. I have purposely refrained, with one exception, from naming individuals, because to do so would require much more space than I am allotted.

I must, however, gratefully acknowledge the continual and professional assistance of Debby Brandt and her staff. Now the job is in the capable hands of **Dale Compton,** who will make his own contributions as he has in so many other roles in the Academy. I wish him well. I will do what I can to help the Academy and to participate in its many important and interesting activities.

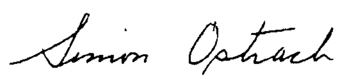

Simon Ostrach
Home Secretary

Third German-American Frontiers of Engineering

The third German-American Frontiers of Engineering Symposium was held at the Fährhaus Meyer-Farge in Bremen, Germany, 13–15 April. Modeled on the U.S. Frontiers of Engineering Symposium, this bilateral Frontiers meeting brought together 60 engineers ages 30–45 from German and U.S. companies, universities, and government agencies. Observers from the United Kingdom, Switzerland, Finland, and the Czech Republic also attended the meeting.

Like the U.S. Frontiers symposium, the goal of the meeting is to bring together emerging engineering leaders in a forum where they can learn about leading-edge developments in a range of engineering fields, thereby facilitating an interdisciplinary transfer of knowledge and methodology. In the case of the bilateral Frontiers, there is the added dimension of helping build cooperative networks of younger engineers that cross national boundaries.

NAE member **Matthew V. Tirrell,** Richard A. Auhil Professor and dean, College of Engineering, University of California, Santa Barbara, and Eduard Reithmeier, head, Institute for Measurement and Control, University of Hannover, co-chaired the organizing committee and the symposium.

The four topics covered at the meeting were bioinformatics, the Internet, production automation, and nanotechnology. Presentations, given by two Germans and two Americans in each of the four areas, covered such topics as bioinformatics and metabolic engineering, future intelligent vehicles and telematic services, trends in agile manufacturing, and biomimetic self-cleaning surfaces. There was interesting and spirited discussion among the participants, both during the formal sessions and during breaks, receptions, and dinners. A highlight of the meeting was a dinner cruise on the Weser River.

Funding for the three German-American Frontiers of Engineering meetings held to date has been provided by the German-American Academic Council Foundation (GAAC) and NAE Foundation funds. With the closure of GAAC operations, the Alexander von Humboldt Foundation will assume sponsorship of these meetings. Plans for a fourth German-American meeting, to be held either in Germany or the United States, are under way.

In tandem with the German-American Frontiers of Engineering symposia, the NAE continues to work on developing a European-American Frontiers of Engineering Symposium that would be carried out in cooperation with European members of the Council of Academies of Engineering and Technological Sciences (CAETS).

For more information about this activity, contact Janet Hunziker in the NAE Program Office at 202–334–1571 or by e-mail at jhunzike@nae.edu.

Council and Staff Awards Presented

The National Academy of Engineering recently hosted its annual Council and Staff Awards Luncheon. Four staff members received NAE staff awards for outstanding service—Robin Gibbin, Karla Weeks, and Maribeth Keitz were honored for their contributions to the successful "Greatest Engineering Achievements of the 20th Century" project, and Kim Garcia was honored for her management of the NAE Council's administrative responsibilities. The honors were presented by **Wm. A. Wulf** and included a certificate of appreciation and a cash award.

In addition to the staff awards, NAE service awards went to Robin Gibbin, Cynthia McFerson, and Greg Pearson for 5 years of service, and to Janet Hunziker and Mary Resch for 10 years of service.

Special recognition went to retiring councillors **Jim Duderstadt, Tom Falkie,** and **Bill Schowalter** for their dedication to serving the NAE and its Council. Their term ended 30 June. President Wulf also acknowledged the dedication and service of Home Secretary **Si Ostrach,** who also retired from NAE service in June.

NAE Calendar of Meetings

2000

7 September	NAE/ABET Leadership Meeting
14–16 September	Sixth Annual Frontiers of Engineering Symposium Irvine, Calif.
18 September	Committee on Engineering Education Chicago, Ill.
25 September	NAE Nominating Committee Russ Prize Committee
25–26 September	Megacities Workshop
27 September	CWE Website Subcommittee
28 September	NAE/AAES IntAC
29 September	NAE/AAES Forum
3 October	NAE/ASEE Engineering Deans Council
10–13 October	International Conference of the Chinese Academy Beijing, China
13 October	CAETS Governing Board Beijing, China
19 October	Great Achievements Book Committee
20–21 October	NAE Council
21–24 October	NAE Annual Meeting
2–4 November	First Japan-America Frontiers of Engineering Nara, Japan
10 November	NAE/Explorers in Engineering & Technology
14–16 November	Forum on Diversity in the Engineering Workforce Dallas, Texas
9–10 December	NAE Committee on Membership Irvine, Calif.

All meetings are held in the National Academies Building, Washington, D.C., unless otherwise noted.

Convocation of Professional Engineering Societies

On 8–9 May the annual convocation of professional engineering societies and the National Academy of Engineering was held in Washington. The convocation is an opportunity for presidents, presidents-elect, and executive directors of the societies and NAE officers to gather and discuss areas of mutual interest. Emphasis is placed on discussion among all participants during the meeting, and all are encouraged to share their experiences, concerns, and questions.

On Monday, 8 May, after a welcome by **Wm. A. Wulf,** NAE president, **John H. Gibbons,** NAE senior fellow, spoke on challenges facing the U.S. engineering community. In the next session, **Tom Young,** retired executive vice president of Lockheed Martin, discussed the topic of technological literacy. He was followed by William E. Dugger, director of the ITEA Technology for All Americans Project, who spoke about standards for the study of technology. **Delores M. Etter,** deputy under secretary of defense, began the afternoon session by addressing emerging issues in the Department of Defense, and then Joseph Bordogna, deputy director of the National Science Foundation, discussed emerging trends in science and technology. The day's program concluded with a panel discussion of public understanding of engineering chaired by Dr. Wulf.

"Diversity and Professional Engineering Community," a panel discussion chaired by Dundee Holt, vice president of the National Action Council for Minorities in Engineering (NACME), began the half-day session on 9 May. The convocation concluded with a general discussion of issues of importance to the U.S. engineering community.

NAE Welcomes Hollomon Fellow

The NAE is pleased to announce that Gretchen E. Matthern is serving a one-year J. Herbert Hollomon Fellowship to assist in the Academy's initiation of the Earth Systems Engineering program. Matthern is an advisory engineer at the Idaho National Engineering and Environmental Laboratory (INEEL) where she facilitates the identification, development, and demonstration of innovative technologies to meet the needs of the INEEL environmental restoration program. Previously she conducted research in microbe-based treatments of hazardous and radionuclide-contaminated wastes and served as the project manager for several environmental remediation actions.

Dr. Matthern earned a B.S. and M.S. in chemical engineering from the University of Oklahoma and a Ph.D. in chemical engineering from the University of Virginia. She joined INEEL in 1988.

NAE Award Recipients Named

Recipients have been named for this year's Arthur M. Bueche and Founders awards. The 2000 Arthur M. Bueche Award goes to **Charles M. Vest,** president, Massachusetts Institute of Technology, and the 2000 Founders Award goes to **Charles H. Townes,** professor, University of California, Berkeley. Please join us at the Annual Meeting on 22 October to honor the recipients!

In Memoriam

JOSEPH BURKE, 86, retired manager, GE Corporate Research and Development Center, died on 29 February 2000. Dr. Burke was elected to the NAE in 1976 for contributions and administration of research and development in ceramics.

HENDRIK B. G. CASIMIR, 90, emeritus professor, Leiden University, died on 4 May 2000. Dr. Casimir was elected a foreign associate of the NAE in 1976 for leadership in research and development of electron tubes, solid-state devices, and glass and metal products.

JAMES A. CAYWOOD III, 77, chairman emeritus, DeLeuw, Cather & Company, died on 10 July 2000. Mr. Caywood was elected to the NAE in 1994 for contributions to the professional and managerial aspects of civil engineering, particularly in the field of transportation.

ALTON C. DICKIESON, 94, retired vice president, Transmission Systems Development, AT&T Bell Laboratories, died on 13 April 2000. Mr. Dickieson was elected to the NAE in 1970 for his role as director of development of long-distance transmission facilities of the Bell system.

JOHN M. HEDGEPETH, 73, president, Digisim Corporation, died on 25 April 2000. Dr. Hedgepeth was elected to the NAE in 1994 for the conception, design, and analysis of structural systems for space science and applications.

S. W. HERWALD, 81, retired vice president, Westinghouse Electric Corporation, died on 20 August 1998. Dr. Herwald was elected to the NAE in 1967 for theory and development of servomechanism systems.

JOHN A. HRONES, 88, retired provost emeritus, Case Western Reserve University, died on 14 June 2000. Dr. Hrones was elected to the NAE in 1975 for contributions as a teacher, administrator, pioneer in the field of automatic control, and leader in engineering education.

I. BIRGER JOHNSON, 87, retired private engineering consultant, died on 1 June 2000. Mr. Johnson was elected to the NAE in 1980 for contributions to the reliable and economic design of power transmission systems and the fundamental understanding of system transient phenomena.

JOSEPH F. KEITHLEY, 84, founder and former chairman, Keithley Instruments, Inc., died on 1 October 1999. Mr. Keithley was elected to the NAE in 1992 for pioneering contributions to electronic test and measurement instrumentation.

W. DAVID KINGERY, 74, Regents Professor, University of Arizona, died on 29 June 2000. Dr. Kingery was elected to the NAE in 1975 for leadership in the science and engineering of ceramic materials, spanning the whole spectrum of physical phenomena, structure-property relationships, innovative processing, and applications to modern technologies.

FRANK W. MCBEE, JR., 80, retired chairman and CEO, Tracor, Inc., died on 7 April 2000. Mr. McBee was elected to the NAE in 1989 for outstanding engineering management leadership resulting in the creation of a unique, world-recognized, high-technology business and for recognized contributions to engineering education.

FRANCIS K. MCCUNE, 94, retired vice president, engineering, General Electric Company, died on 10 May 2000. Mr. McCune was elected to the NAE in 1966 for contributions to nuclear engineering.

E. R. PIORE, 91, retired vice president, director, and chief scientist, IBM Corporation, died on 9 May 2000. Dr. Piore was elected to the NAE in 1966 for research policy and planning in electronics and solid-state technology.

National Research Council Update

Library of Congress Faces Digital Information Challenge

The Library of Congress must act quickly to address strategy, management, funding, and staffing issues in order to face the challenges posed by the explosion of digital information. This is the conclusion of a new National Research Council report, *LC21: A Digital Strategy for the Library of Congress,* prepared by a study committee of the NRC's Computer Science and Telecommunications Board.

To maintain its important role in the political, intellectual, and cultural life of the nation, the library must continue to handle traditional media while making the transition to collecting and preserving electronic books, news websites, digital music, and other forms of electronic information.

According to the report, the library must develop robust strategies to handle this enormous task, including strategies for improving organizational capabilities and resources. To that end, the report recommends that the library create a new deputy librarian position for strategic initiatives; establish a group to guide the library's information technology strategies and efforts; and establish a technical advisory board with members drawn from outside the library.

A key limitation of the library is its information infrastructure, which will require significant upgrades to transition to collecting digital objects. In particular, the current infrastructure is lacking in networking and security capabilities. Overall, the library needs a new system for digital objects that is integrated with the established system for acquiring and archiving physical formats, including provisions for registering and depositing digital objects with the U.S. Copyright Office (a unit of the library).

Other challenges the library faces include the legal, economic, and technical issues associated with preserving digital information. The library should lead efforts to address these issues, but should also initiate cooperative work with electronic publishers and the R&D community, and establish contractual arrangements with publishers and distributors of significant digital content.

The report was sponsored by the Library of Congress. NAE members **James Gray** and **Jerome Saltzer** served on the study committee.

continued from page 34

SIGURD A. SJOBERG, 80, retired deputy director, NASA Johnson Space Flight Center, died on 26 March 2000. Dr. Sjoberg was elected to the NAE in 1974 for contributions to the advancement of aeronautics and manned space-flight technology.

JUDSON S. SWEARINGEN, 92, retired president, Rotoflow Corporation, and consultant, died on 5 September 1999. Dr. Swearingen was elected to the NAE in 1977 for contributions to the technology of expanders and compressors, to cryogenic systems, and to the design of shaft seals for high-speed machinery.

EWALD WICKE, 85, professor emeritus, University of Muenster, died on 7 March 2000. Dr. Wicke was elected a foreign associate of the NAE in 1983 for leadership in the development of chemical engineering education and research in Germany and for contributions to chemical reaction engineering.

Publications of Interest

The following publications result from the program activities of the National Academy of Engineering or the National Research Council. Except where noted, each publication is for sale (prepaid) from the National Academy Press (NAP), 2101 Constitution Avenue, N.W., Lockbox 285, Washington, DC 20055. For more information or to place an order, contact NAP online at http://www.nap.edu or by phone at (800) 624–6242. *(Note: Prices quoted by NAP are subject to change without notice. Online orders receive a 20 percent discount. Please add $4.50 for shipping and handling for the first book ordered and $0.95 for each additional book. Add applicable sales tax or GST if you live in CA, DC, FL, MD, MO, TX, or Canada.)*

Biobased Industrial Products: Research and Commercialization Priorities. Discusses the concept of the biorefinery and outlines proven and potential thermal, mechanical, and chemical technologies for conversion of natural resources to industrial applications; illustrates the developmental dynamics of biobased products through existing examples, as well as products still on the drawing board; and identifies priorities for research and development. Paperbound, $35.00.

The Digital Dilemma: Intellectual Property in the Information Age. Presents the multiple facets of digitized intellectual property, defining terms, identifying key issues, and exploring alternatives. Follows the complex threads of law, business, commercial incentives, and information rights, and proposes research and policy recommendations as well as principles for policymaking related to intellectual property rights. Paperbound, $42.95.

Frontiers of Engineering: Reports on Leading Edge Engineering from the 1999 NAE Symposium on Frontiers of Engineering. Presents extended summaries of the presentations given at the October 1999 symposium on topics that include the human genome, engineering novel structures, and energy for the future and its environmental impact. Paperbound, $31.00.

LC21: A Digital Strategy for the Library of Congress. Examines the impact of digital information on the Library of Congress and makes recommendations for developing a new system for digital objects that is integrated with the existing systems for acquiring and archiving physical formats. Paperbound, $37.00.

Making IT Better: Expanding Information Technology Research to Meet Society's Needs. Highlights the fundamental importance of research to ensure that information technology, a critical underpinning to our nation's success, keeps pace with society's expanding needs. Forthcoming. Paperbound, $29.95.

Renewable Power Pathways: A Review of the U.S. Department of Energy's Renewable Energy Programs. Reviews the DOE Office of Power Technologies (OPT) and makes recommendations for OPT as a whole as well as for individual OPT programs, including research into biopower, hydropower, geothermal, solar, and hydrogen energy generation methods. Paperbound, $30.75.

Review of the Research Program of the Partnership for a New Generation of Vehicles: Sixth Report. Examines the overall adequacy and balance of the PNGV research program. Discusses ongoing research on fuels, propulsion engines, and emission controls and reviews the USCAR partners' progress on PNGV concept vehicles for 2000. Paperbound, $32.00.

Trust in Cyberspace. Provides an assessment of current state-of-the-art procedures for building trustworthy networked information systems, proposing directions for research in computer and network security, software technology, and system architecture. Assesses current technical and market trends in order to better inform public policy as to where progress is likely and where incentives could help. Hardcover, $44.95.