


- 3 Newslog
- 10 Forum
- 14 Book reviews
- 16 Software review
- 17 The engineer at large
- 18 Innovations
- 19 Program notes
- 22 Video
- 24 Speakout
- 26 Faults & failures
- 58 EEs' tools & toys
- 60 Papers are invited
- 61 Calendar
- 63 New publications
- 67 IEEE tables of contents
- 82 Scanning THE INSTITUTE
- 82 Coming in *Spectrum*

Cover: Before destroying all intermediate-range missiles under the recently signed arms pact, Soviet and United States workers may salvage the nuclear warheads. This uranium-235 button, the end product of a recovery process at the Oak Ridge Y-12 Plant in Tennessee, may be recast and machined for warhead components. The 9.5-lb button, actually lead-colored, is valued at about \$200 000. Photo: Martin Marietta Energy Systems Inc.



Field Editor Perry pays for a tank of gas with her bank credit card

46



Minicomputer pioneer Bell lectures on the merits of a national network

54

STAFF

Editor and Publisher: Donald Christiansen
Executive Editor: Edward A. Torrero
Administrative Editor: Ronald K. Jurgen
Senior Technical Editor: Gadi Kaplan
Managing Editor: Mark A. Fischetti
Senior Associate Editor: Trudy E. Bell
Associate Editors: John A. Adam, Elizabeth Corcoran, Karen Fitzgerald, Erin E. Murphy, Gary Stix, John Voelcker, Katherine Wollard, Glenn Zorpette
Field Editor (West Coast): Tekla S. Perry
Chief Copy Editor: Margaret E. Eastman
Copy Editor: Michael Olver
Editorial Researcher: Robert Davidson
Contributing Editors: Evan Herbert, Alexander A. McKenzie, Michael F. Wolff
Manuscript Review: Rita Holland (Supervisor)
Design Consultant: Gus Sauter
Art and Production: Janet Mannheim (Director)
Editorial Production: Ruth M. Edmiston (Manager), Marcia Meyers (Production Supervisor), Peter Ruffett (Typographer)
Art: Morris Khan (Technical Graphic Artist), Elaine Wesson Da Costa (Manager)
Advertising Director and Associate Publisher: William R. Saunders
Administrative Assistant: Wendy Tinkler
Research: Hendrik Prins (Manager)
Circulation and Research Administrator: Joanne Biley
Mail List Sales: Geoffrey L. Klapisch (Manager)
Business Manager: Carl Maier
Marketing: Robert D. Moran (Manager)
Administrative Assistant to the Editor and Publisher: Nancy T. Hantman

ADVISORY BOARD

Chairman: Jerome J. Suran
 Edward C. Bertnoff, Richard W. Damon, N. Rex Dixon, Thelma Estrin, Ronald G. Hoelzeman, Robert W. Lucky, G.P. Rodrigue, Michael J. Whitelaw


EDITORIAL BOARD

Chairman: Donald Christiansen
 Wallace B. Behnke, Peter D. Blair, James E. Brittain, S.J. Campanella, Walter S. Ciora, Theo A.C.M. Claasen, Almon H. Clegg, Karsten E. Drangeid, Lester Eastman, Kimberly D. Emerson, Ernst Feldtkeller, Bob L. Gregory, Frederick Hayes-Roth, Charles K. Kao, Fred L. Katzmann, John L. McLucas, M. Granger Morgan, Eric Nussbaum, Alfred R. Potvin, Howard Sachs, John B. Slaughter, Cary Spitzer, Michiyuki Uenohara, Thomas P. Woll

Postal identification statement is on page 6.

Gordon Bell calls for a U.S. research network

Advances in computing and communications research in the United States will depend on a powerful data-communications network, perhaps best provided by the Federal government



If a research cardiologist at Boston University Medical Center urgently needs to review cardiac images with a colleague at the Mayo Clinic in Minnesota, he can either express-mail his material to the clinic or fly there. What he cannot do is transmit the material instantaneously from the computer workstation in his office. On the other hand, researchers at the Massachusetts Institute of Technology do not

lack networks. They can communicate with many research organizations around the world, but they must use the right one of a dozen networks to do it.

These scenarios point up just two absurdities of the present situation in U.S. computer networking. Existing networks not only lag well behind the growing needs of the research community—they are too fragmented to develop unaided into a single, coherent system.

The most viable solution is a national research network organized and maintained by the Federal government. Access to information has never been more important than it is today, and the ability to fully exploit information resources—be they individual researchers, research teams, databases or supercomputers—will determine how competitive any group or nation is. Any new proposal costs, of course. But a single national network, jointly supported by all the Government agencies now running independent networks, could well save money.


Quantity does not equal quality

Computer networking among scattered facilities in the United States began in 1969, when the Defense Advanced Research Projects Agency (Darpa) established Arpanet. The network started out as a means of sharing expensive equipment, databases, files, and above all time on computers that would otherwise have lain idle. What developed, however, was a completely different style of interaction. Utilizing the ability to send mail and large documents electronically, researchers have built electronic bulletin boards and held extensive forums and conferences.

Today's Arpanet is conceptually identical to the network of the early 1970s. But it can do little more than swap computer-mail messages now that the number of machines has mushroomed beyond a hundred switching computers that connect hundreds more shared computers and workstations. The network could be upgraded, with great difficulty. But the Defense Communications Agency, which oversees it, is reluctant to run a civilian research network.

Since Arpanet was established, some 36 other major research networks have sprung up around the world, all based on variations of Arpanet's method for packet-switching data. Typically

C. Gordon Bell Ardent Computer Corp.



If science is willing to wait, a national network just might eventually evolve, according to Bell. But patience seldom assists progress.

Some are addicted to information as predigested experience, a cognitive fast food. This is a danger for the new generation, raised on the instant knowledge of television and the computer. They need discipline to filter and edit useful information, and beyond this, to develop deeper interest and understanding of the world and of themselves.

Some self-developers feel their possibilities are unlimited and so are unrealistic about what they can achieve. They seem like neglected, hyperactive children as they flit compulsively from one "developmental experience" to another.

Commitment and intimacy are problems for self-developers. "How much of myself am I willing to commit?" is a popular phrase for this new generation. Of the self-developers we interviewed, a high percentage (25 percent) have been divorced and fewer than half are in their first marriage (48 percent). In contrast, only 6 percent ever divorced in *The Gamesman* sample of high-tech managers. Given that self-developers are on the average younger than the other types, the difference may become even greater in the future, as more self-developer marriages break up.

Motivating the self-developer

To motivate the new generation, managers must give them opportunities to develop marketable skills. Self-developers are also stirred by the chance to expand their knowledge, improve their well-being, live a more enriching life through travel abroad.

They see money as only part of a job's reward and, as long as their basic financial needs are met, weigh income against other payoffs like the opportunity to develop skills, time off, health care, child care, exercise facilities, and a friendly atmosphere.

When managers fail to make work meaningful to self-developers, the self-developers will find ways to express themselves outside the company—or will quit. In fact, they switch jobs more easily than traditional professionals. This willingness to quit an oppressive job, the dual-career family, money in the bank, and secondary entrepreneurial ventures, all help to fan the spirit of independence.

Self-developers find it relatively easy to jump ship because, they say, work isn't the only important thing in their lives. The ones with families want to keep a balance. A highly competent and motivated executive, aged 38, says: "I work 50 hours a week. I come in early and go home late, but I leave it at the office. I am not going to push my little son off my lap because he is messing up papers from work."

For self-developers, physical well-being also requires attention. A 29-year-old manager of information systems says: "We watch the things we eat. We run. If you're a workaholic, after a while you're not really productive."

The move into management

While self-developers are critical of the managerial hierarchy, they will move up, albeit reluctantly: it is safer and crisper to be a professional self-developer than to commit oneself to risky projects and the education of others. But they recognize they need power to get things done.

The good news is that self-developers are natural facilitators. They can create an open atmosphere where views are exchanged, conflict becomes constructive dialogue and study, and consensus is achieved. They can facilitate well because they are egalitarian and interested in other's views and ideas.

All self-developer managers believe that to succeed they must create a motivated team. For them, being a team player does not mean group thinking. It means playing a special role on a team where each player has a say in how to implement strategy.

Self-developers succeed best as managers when they institute good practices: frequent evaluations, team meetings, and training in group process and problem-solving.

The weaknesses of self-developers as managers stems from their reluctance to make commitments, which appears to me a significant problem in about 40 to 50 percent. And no one gains authority in the minds of others without commitment to projects and

The tellers of the tale

The material in this article is excerpted from *Why Work*, to be published in March by Simon & Schuster, New York. While my earlier book, *The Gamesman*, focused on the elite, high-technology company, the study upon which *Why Work* is based was much broader. Its conclusions apply to engineers as well as to other professionals in today's technoservice economy.

Why Work is the result of seven years of research involving nine companies and more than a dozen Federal, state, and municipal agencies. Over 350 people at all levels, from chief executives to front-line service employees, were interviewed about their work values in sessions lasting from 1½ to 3 hours. They were asked what satisfied and dissatisfied them at work, how they defined service, how they wanted to be managed, how they managed others, how they relate to customers, clients, and co-workers, and about their family background and goals at work. Many of the questions were similar to those used by *Spectrum* in its series of articles by Tekla S. Perry from December 1983 to May 1984 on engineering environments.

Businesses where employees were interviewed and surveyed include AT&T, U.S. West, a large big-eight accounting firm, an innovative insurance company, a TV broadcasting company, a large supermarket chain, the service division of a large oil company, a company producing information systems, and Scandinavian Airline Systems. The government agencies include the Internal Revenue Service; the Commerce Department; Action (which comprises the Peace Corps, Vista, and the Older Americans Volunteer Program); the Federal Aviation Administration; the National Aeronautics and Space Administration; the departments of agriculture, justice, and defense; the Federal Trade Commission; the Veterans Administration; the National Highway Traffic Commission; the Library of Congress; a statewide health department; two hospitals; a county health department; a city tax office; a social worker office in California; and a municipal library. Besides these, the study drew on interviews conducted by my colleagues and students with police officers in two metropolitan departments, entrepreneurs in the U.S. and Sweden, and middle managers in Japanese banks and trading companies.

—M.M.

people. If self-developers are going to become effective leaders, they must find meaning in caring for others and taking responsibility for larger enterprises.

To probe further

The motivation of entrepreneurs and managers in high-technology industry is described in Michael Maccoby's best-seller, *The Gamesman* (Simon & Schuster, New York, 1977). This book originated as a *Spectrum* article, "Winning and losing at work" [July 1973, pp. 39-48].

A new model of leadership for the 1980s that combined improved competitiveness and the quality of working life is described in Maccoby's book, *The Leader* (Simon & Schuster, 1981).

The relationship of employees of all levels in major corporations and government agencies to their workplace is analyzed in *Why Work*, a book by Maccoby to be published in March by Simon & Schuster.

About the author

Michael Maccoby directs the Project on Technology, Work and Character, a center for research and consulting in Washington, D.C., and he also acts as consultant to business, government, and labor unions. He has a Ph.D. in social relations (1960) from Harvard University. ◆

these networks convey data at rates of 1.2 to 56 kilobits per second, using incompatible communication protocols.

Federal agencies usually support an average of two independent wide-area networks. Often these networks go to different buildings on the same site—be it university campus or Federal laboratory—wasting resources. Yet the Government cannot even begin to estimate the current costs because each Federal agency considers its networking expenses proprietary information.

In any event, more networks do not automatically translate into greater capabilities. The situation is reminiscent of telephone systems in the early 1900s, when a town might support several distinct company telephone networks, forcing subscribers to use a deskful of phones. Theodore N. Vail, president of the American Telephone & Telegraph Co., however, successfully corralled the local companies under the banner: "One policy, one system, universal service." Similarly today, easily a dozen incompatible networks may overrun just one university campus.

Their incompatibility, if not their numbers, is fortunately waning. Within the past two years, most networks have begun migrating toward Darpa's Transmission Control Protocol/Internet Protocol (TCP/IP) standards and are committed to using the internationally approved Open Systems Interconnect (OSI) standards as they become available. Identical protocols for exchanging data not only make it easy to connect networks but to have them share common links to save equipment costs.

On university campuses, researchers are wiring their array of local-area networks (LANs) into campus area networks (CANs). But at present, most campus area networks are isolated. Since LANs typically operate at 10 megabits per second and CANs operate at up to 100 Mbits/s, very fast wide-area networks will be required to connect these CANs into a global network.

Faster computers, fiercer demands

The rise in computer speed on campuses has triggered other developments. It has revolutionized the nature of the data that researchers want to share, boosted the demands on supercomputers, and encouraged new forms of collaborative research.

A report released last July from the National Science Foundation (NSF) stressed how sorely researchers need networking to create visual and hence more comprehensible representations of supercomputer output and to exchange high-quality graphical data, including photographs. Since a high-resolution workstation displays about a million pixels that each can change as often as 10 to 60 times a second, the bandwidth required for sending dynamic pictures varies from 10 to 60 Mbits/s for a black and white display. For a color display, it soars to 320 to 1920 Mbits/s.

These predictions assume data is not compressed before being transmitted. Depending on the application, data compression techniques can reduce the necessary bandwidth by a factor of 10 to 1000. Nevertheless, the data requirements of a national research network still exceed current capabilities.

Connecting several hundred workstations and high-performance computers would require a network capable of delivering hundreds of megabits per second. Visualizing mechanical parts, medical images, and geological data can demand the transfer of some 4 gigabytes of data among workstations or from supercomputers to workstations. Without data encoding, a 45-Mbit/s link would transmit those 4 Gbytes within 10 minutes. Today's networks are over 1000 times slower!

Remote access to supercomputers is particularly critical since few institutions can afford their own. As supercomputing time is limited, researchers usually run pieces of a program on smaller computers or workstations, then transmit the entire program and database over a network to a supercomputer. Bandwidth requirements for such transmissions easily exceed 1 Mbit/s.

Increasingly, the problems under study require the active collaboration of researchers who are scattered among various research institutions. Collaboration technology—an emerging area in which researchers work together over a network—depends on a range of networking capabilities, including: compound docu-

ment transfer and the simultaneous viewing and editing of documents that combine text, graphics, pictures, and voice; computer conferencing with attendees interacting through both pictures and conversation; design reviews performed on a common document; and the ability to remotely control and interact with special laboratory and industrial facilities.

Clearly, huge volumes of data will pour through the national research network. Other countries are already acting: the Japanese and European governments are busy building fiber-optic computer networks that will transmit gigabits of data per second. In Japan, most major research centers already plug into high-speed networks that enable them to store and transmit international scientific and engineering material.

Spontaneous generation unlikely

All these developments point to the need for a national research network in the United States. But might that network evolve over time, without a centrally administered plan?

Some argue that a sort of national network may emerge with advances in fiber-optic technology. Indeed, today's fiber-optic communication links promise 1000 to 100 000 times the capacity and speed of traditional cable and satellite channels. But the cost of creating a large enough fiber-optic network remains prohibitive.

Supply and demand play no part in the price of fiber-optic networks. A vast amount of fiber remains unused. But the prices of using fiber-optic links are primarily based on the rates for transmitting voice communication on coaxial cables. Since a fiber-optic bundle can transmit several orders of magnitude more data than a coaxial cable, the price per bit transmitted should be quite low. Carriers may blame regulatory agencies for the voice-based pricing, but the situation almost forces any organization wishing to send data in this way to set up a private network.

Furthermore, switching equipment to exploit even the DS3 standard of 45 Mbits/s is unavailable. Researchers have developed some designs and prototypes for 45-Mbit/s packet switches. But no companies currently manufacture such products. Suppliers of communications services are loath to invest in fast data-communications networks, since they already have large voice-traffic networks and many customers. Proposals for national standards for higher speeds are moving at glacial rates.

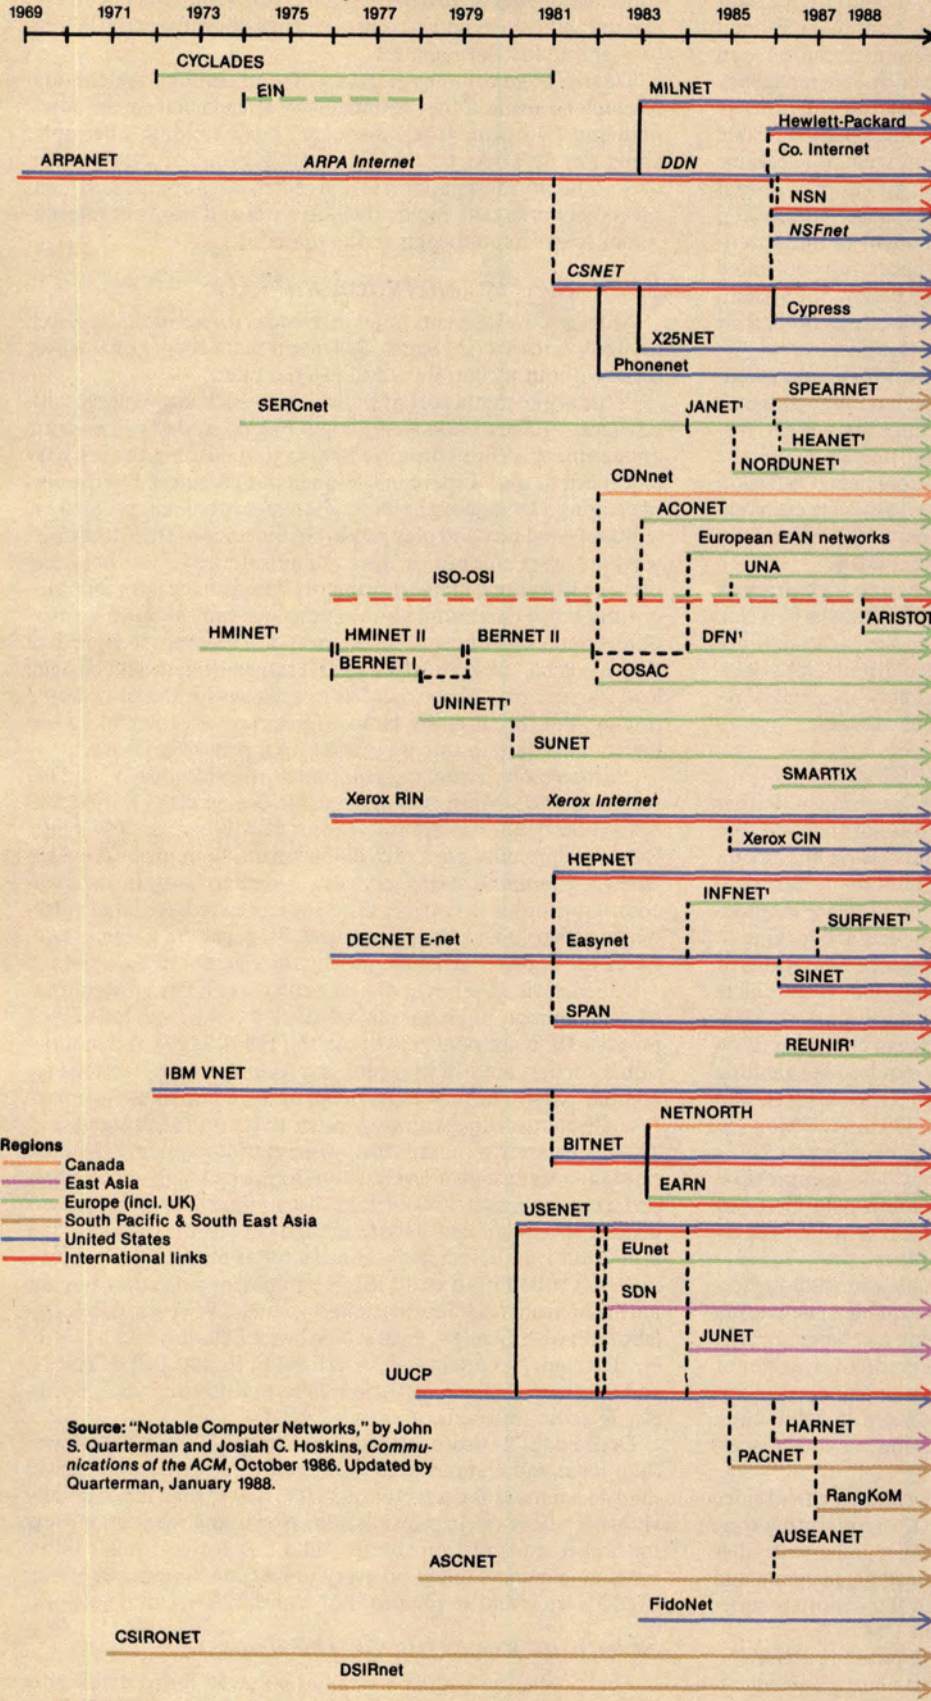
The integrated-services digital network (ISDN) is often touted as a panacea. But it has moved much too slowly to hold much promise for a national network in the 1990s. Local and international carriers are still thrashing out technical specifications for compatibility. The regional Bell operating companies are not especially cooperating with each other to set up ISDN standards. There are barely standards for low-level protocols, and there are no standards at higher levels. Moreover, when high-speed fibers do terminate in switching offices, distributing them to local users takes an inordinate amount of time and effort—the so-called "last mile" problem. In fact, U.S. manufacturers are losing the ability, if they ever had it, to build ISDN equipment since they buy the hardware from their international partners. Wisconsin Bell collaborates with Siemens. Pacific Northwest Bell works with Northern Telecom. Mountain Bell, AT&T, GTE, Illinois Bell, NYNEX, and Southern Bell all have various links to Ericsson, NEC, Northern Telecom, Siemens, and Fujitsu.

Defined ISDN protocols remain a factor of 200 to 2000 slower than local and campus area networks, a factor of 1000 slower than 45 Mbits/s, and a factor of 30 000 slower than is potentially viable for fiber-optic transmission. Nor is any work under way for higher transmission speeds. Yet a link for 45-Mbit/s transmission is needed today on every university campus. In effect, ISDN is irrelevant to the needs of a national research network.

Solutions knocking on the door

The closest the United States has come so far to developing a national research network is the 56-kbit/s NSFNET. Last November, the NSF began to extend the NSFNET backbone beyond the present six supercomputing centers to include seven

Time lines of notable computer networks



- ACONET: Akademisches Computer Netz
- ARISTOTE: Association de Réseaux Informatiques en Système Totalement Ouvert et Très Elaboré
- ARPANET: Advanced Research Projects Agency Network
- ASCNET: Australian Computer Science Network
- AUSEANET: Austroasian Network
- BERNET I, II: Berlin Network
- BITNET: Because It's Time Network
- COSAC: Communications sans Connexions
- CSIRONET: Commonwealth Scientific & Industrial Research Organization Network
- CSNET: Computer Science Network
- DDN: Defense Data Network
- DECNET E-net: DEC Engineering Network
- DFN: Deutsche Forschungsnetz
- DSIRnet: Government network in New Zealand
- EARN: European Academic Research Network
- EIN: European Informatics Network
- EUnet: European Unix Network
- HARNET: Hong Kong Academic & Research Network
- HEANET: Higher Education Authority Network
- HEPNET: High Energy Physics Network
- HMINET I, II: Hahn-Meitner Institut
- IBM VNET: Virtual Network
- INFNET: Istituto Nazionale Fisica Nucleare
- ISO-OSI: International Organization for Standardization-Open Systems Interconnect
- JANET: Joint Academic Network
- JUNET: Japanese Unix Network
- MILNET: Military Network
- NORDUNET: Nordic University Network
- NSFnet: National Science Foundation Network
- NSN: NASA Science Network
- PACNET: Pacific Network
- RangKoM: Rangkaian Komputer Malaysia
- REUNIR: Réseaux des Universités et de la Recherche
- SDN: System Development Network
- SERCnet: Science Engineering Research Council Network
- SINET: Schlumberger Information Network
- SPAN: Space Physics Analysis Network
- SPEARNET: South Pacific Education and Research Network
- SUNET: Swedish University Network
- SURFNET: Dutch university network
- UNA: Universitäts-Netz Austria
- UNINETT: Nordic university network
- USENET: Users' Network
- UUCP: Unix to Unix Copy
- Xerox CIN: Corporate Internet
- Xerox RIN: Research Internet

Note: CDNnet, CYCLADES, EAN, NETNORTH, and SMARTIX are not acronyms.

- Regions**
- Canada
 - East Asia
 - Europe (incl. UK)
 - South Pacific & South East Asia
 - United States
 - International links

Source: "Notable Computer Networks," by John S. Quarterman and Josiah C. Hoskins, *Communications of the ACM*, October 1986. Updated by Quarterman, January 1988.

The past five years have seen the number of networks soar dramatically. Many initially used transmission protocols and technology developed by one or more older networks; these are indicated by the vertical dashed lines connecting networks. (More recently, some networks have begun to use other protocols, particularly ISO-OSI and Arpanet standards.) Solid vertical lines between networks indicate systems under closely related administrations. Dashed horizontal lines indicate protocols or demonstration systems, rather than operational networks. Networks in italics are internets—several networks tied together that use the same transmission protocols.

¹Participants in the Réseaux Associés pour la Recherche Européenne network (RARE)

regional, university-based research networks. The backbone should be running at 1.5 Mb/s by July under the management of Merit Inc., which is based at the University of Michigan, Ann Arbor, and assisted by IBM Corp. and MCI Communications Corp. Unfortunately, the fiscal 1988 budget allocated by Congress for NSFNET barely keeps the network alive.

But the NSF is only one agency. It has no authority for incorporating other Federal networks, no timetable for upgrading to 45-Mbit/s rates, and certainly no budget for doing so.

The U.S. government is slowly recognizing the need for a national research network. In 1986, Congress requested that the Office of Science and Technology Policy (OSTP) study the problems and options of developing a communications network for research computers, including supercomputers at U.S. universities and Federal research facilities, and provide a plan for action by August 1987. The OSTP accordingly established a new inter-agency group, the Federal Coordinating Council for Science, Engineering, and Technology for Computer Research and Applications. The council finished a three-volume report in on time; the OSTP finally sent a summary of the report to Congress late last November.

I chaired the subcommittee on computer networking, infrastructure, and digital communications. In the report, we strongly urged that the Government create a national research network to "foster and enhance the U.S. position of world leadership in computer networking." I believe the situation is far worse; we have already lost leadership in this field. By developing a network that enables U.S. researchers at all universities, national labs, and companies to share resources and ideas, the country just might regain its footing.

Implementing the network can be done in three steps. Stage 1 should be to connect Arpanet with other networks supported by Federal agencies over the next two years. If coordinated and centrally managed, these facilities could unite many computer networks into a seemingly single computer network. Operating the backbone and major regional networks at 1.5 Mb/s should open up a whole new set of library and educational services.

The Government should provide funds for stage 1. The annual cost for such an upgraded service is likely to be \$5 million and should be shared by the five Federal agencies that support the most networking: NSF, Darpa, the Department of Energy, the National Aeronautics and Space Administration, and the Department of Health and Human Services. As part of stage 1, a network manager should be selected and made responsible for upgrading the network speed from 1.5 to 45 Mb/s within three years (stage 2) and to many gigabits a second by the late 1990s (stage 3).

Stage 2 should include upgrading and expanding the existing facilities at 200 to 400 U.S. research institutions at data communication rates of 1.5 Mb/s, or T1 rates. This work would require new funding at approximately \$5 million per year over five years. The estimate assumes that the price of T1 lines will halve over the next five years—a modest assumption, since oversupply pushed prices down more than that in the second half of 1987 alone. Operating expenses for the upgraded facilities are likely to be \$50 million annually. While Government should support the first years, eventually users should cover the costs of the network service, the same way they now pay for telephone service.

Establishing a vigorous, focused program of network research and development is critical to stage 3. Some \$400 million would be needed over 10 years to advance networking technology and make it possible to transmit and switch 3 Gb/s by the early 2000s. Such a network would have 100 000 times more capacity than those currently available and enable researchers to communicate instantaneously.

Who will take the lead?

The Government is certainly not the only hope in this situation. Any one or a combination of the existing telecommunications suppliers could pre-empt Federal efforts to build a nation-

al research network by simply building the network and offering the service for sale. A highly aggressive, imaginative telecommunications industry could view the network as the major, large-scale social experiment of this century and far-sighted preparation for the next.

Achieving stage 2 (45 Mb/s) in three years would take only a small fraction of this industry's research, development, and operations budget. Government should encourage such efforts in every way possible, but, to judge from history, the country cannot depend on such an initiative.

The OSTP report took a different tack. It recommended appointing a lead agency to oversee networking. In January, NSF volunteered for the job; the agencies that participated in the OSTP report agreed. Still, there are no examples of a single agency supplying a facility for the entire research enterprise. In fact, agency behavior is byzantine: each wants its own facility, no matter what the cost. Both the executive and legislative branches of the Federal government are simply incapable of setting up a line item to be funded and administered either by an agency or by an inter-agency group, even though the facility would support the entire research and higher education community.

A radical approach, though, could work: select a private-sector company to manage and develop the network, and provide it with a budget, to which every agency would contribute under NSF guidance. (Each would list its support for the network as a single line item in its budget. Each would also relinquish control of its networks to the manager.) Assure the network manager of steady support—both fiscal and political—for the first five operating years or so of the network. And instruct the manager to devise a plan for gradually shifting all operating costs to users.

Both the national research network and supercomputer facilities could be funded in this fashion. For lack of a common facilities budget, the Federal agencies at present have no choice but to fund and build their own, inevitably overlapping networks and supercomputer centers. Perpetuating this situation only wastes more dollars and more time.

Building this network is not a difficult problem for the U.S. engineering community. But the United States lacks leadership in communications as well as anything resembling a coordinated Federal science and technology policy. Our best hope may be the research community, if it can successfully mobilize its resources. After all, it stands to benefit the most and the soonest from a national research network, even if the country's strength and prosperity over the long haul is what is ultimately at stake.

To probe further

The summary report, "A Research and Development Strategy for High-Performance Computing," sponsored by the Office of Science and Technology Policy, includes an outline for a national research network. Copies are available from the OSTP, attn. Kathleen Bernard, New Executive Office Building, Room 5005, Washington, D.C. 20506. Thoroughgoing descriptions of many existing research networks are given in "Notable Computer Networks," by John S. Quarterman and Josiah C. Hoskins, *Communications of the ACM*, vol. 29, no. 10, Oct. 1986, pp. 932-968. *IEEE Communications Magazine* has also carried many articles on computer networking. Among them, "Research computer networks and their interconnection," by L.H. Landweber, D.M. Jennings, and I. Fuchs, June 1986, pp. 5-17, is a good introduction.

About the author

Best known for his work on the VAX superminicomputer, C. Gordon Bell [F] is vice president for research and development at Ardent Computer Corp., Sunnyvale, Calif. (formerly Dana Computer). Under the Federal Coordinating Council for Science, Engineering, and Technology, Bell chaired the subcommittee on computer networking, infrastructure and digital communications. In 1986 and 1987, he served as assistant director for computing at the National Science Foundation. He earned his B.S. and M.S. degrees in electrical engineering at MIT. ♦