



ONLY CONNECT: Researcher Hubert Zimmermann [left] explains computer networking to French officials at a meeting in 1974. Zimmermann would later play a key role in the development of the Open Systems Interconnection standards.



THE INTERNET THAT WASN'T

The making—and forgetting—of the
Open Systems Interconnection standards

BY ANDREW L. RUSSELL

If everything had gone according to plan, the Internet as we know it would never have sprung up. That plan, devised 35 years ago, instead would have created a comprehensive set of standards for computer networks called Open Systems Interconnection, or OSI. Its architects were a dedicated group of computer industry representatives in the United Kingdom, France, and the United States who envisioned a complete, open, and multilayered system that would allow users all over the world to exchange data easily and thereby unleash new possibilities for collaboration and commerce.

For a time, their vision seemed like the right one. Thousands of engineers and policy-makers around the world became involved in the effort to establish OSI standards. They soon had the support of everyone who mattered: computer companies, telephone companies, regulators, national governments, international standards setting agencies, academic researchers, even the U.S. Department of Defense. By the mid-1980s the worldwide adoption of OSI appeared inevitable.

And yet, by the early 1990s, the project had all but stalled in the face of a cheap and agile, if less comprehensive, alternative: the Internet's Transmission Control Protocol and Internet Protocol. As OSI faltered, one of the Internet's chief advocates, Einar Stefferud, gleefully pronounced: "OSI is a beautiful dream, and TCP/IP is living it!"

What happened to the "beautiful dream"? While the Internet's triumphant story has been well documented by its designers and the historians they have worked with, OSI has been forgotten by all but a handful of veterans of the Internet-OSI standards wars. To understand why, we need to dive into the early history of computer networking, a time

when the vexing problems of digital convergence and global interconnection were very much on the minds of computer scientists, telecom engineers, policymakers, and industry executives. And to appreciate that history, you'll have to set aside for a few minutes what you already know about the Internet. Try to imagine, if you can, that the Internet never existed.

The story starts in the 1960s. The Berlin Wall was going up. The Free Speech movement was blossoming in Berkeley. U.S. troops were fighting in Vietnam. And digital computer-communication systems were in their infancy and the subject of intense, wide-ranging investigations, with dozens (and soon hundreds) of people in academia, industry, and government pursuing major research programs.

The most promising of these involved a new approach to data communication called packet switching. Invented independently by Paul Baran at the Rand Corp. in the United States and Donald Davies at the National Physical Laboratory in England, packet switching broke messages into discrete blocks, or packets, that could be routed separately across a network's various channels. A computer at the receiving end would reassemble the packets into their original form. Baran and Davies both believed that packet switching could be more robust and efficient than circuit switching, the old technology used in telephone systems that required a dedicated channel for each conversation.

Researchers sponsored by the U.S. Department of Defense's Advanced Research Projects Agency created the first packet-switched network, called the ARPANET, in 1969. Soon other institutions, most notably the computer giant IBM and several of the telephone monopolies in Europe, hatched their own ambitious plans for packet-switched networks. Even as these institutions contemplated the digital convergence of computing and communications, however, they were anxious to protect the revenues generated by their existing businesses. As a result, IBM and the telephone monopolies favored packet switching that relied on "virtual circuits"—a design that mimicked circuit switching's technical and organizational routines.

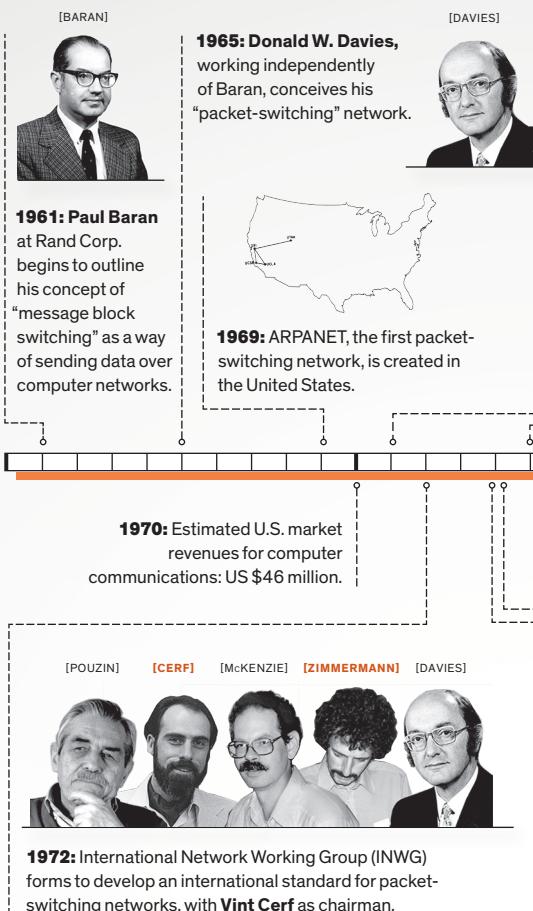
With so many interested parties putting forth ideas, there was widespread agree-

ment that some form of international standardization would be necessary for packet switching to be viable. An early attempt began in 1972, with the formation of the International Network Working Group (INWG). Vint Cerf was its first chairman; other active members included Alex McKenzie in the United States, Donald Davies and Roger Scantlebury in England, and Louis Pouzin and Hubert Zimmermann in France.

The purpose of INWG was to promote the "datagram" style of packet switching that Pouzin had designed. As he explained to me when we met in Paris in 2012, "The essence of datagram is connectionless. That means you have no relationship established between sender and receiver. Things just go separately, one by one, like photons." It was a radical proposal, especially when compared to the connection-oriented virtual circuits favored by IBM and the telecom engineers.

INWG met regularly and exchanged technical papers in an effort to reconcile its designs for datagram networks, in particular for a transport protocol—the key mechanism for exchanging packets across different types of networks. After several years of debate and discussion, the group finally reached an agreement in 1975, and Cerf and Pouzin submitted their protocol to the international body responsible for overseeing telecommunication standards, the International Telegraph and Telephone Consultative Committee (known by its French acronym, CCITT).

The committee, dominated by telecom engineers, rejected the INWG's proposal as too risky and untested. Cerf and his colleagues were bitterly disappointed. Pouzin, the combative leader of Cyclades, France's own packet-switching research project, sarcastically noted that members of the CCITT "do not object to packet switching, as long as it looks just like circuit switching." And when Pouzin complained at major conferences about the "arm-twisting" tactics of "national monopolies," everyone knew he was referring to the French telecom authority. French bureaucrats did not appreciate their countryman's candor, and government funding was



drained from Cyclades between 1975 and 1978, when Pouzin's involvement also ended.

For his part, Cerf was so discouraged by his international adventures in standards making that he resigned his position as INWG chair in late 1975. He also quit the faculty at Stanford and accepted an offer to work with Bob Kahn at ARPA. Cerf and Kahn had already drawn on Pouzin's datagram design and published the details of their "transmission control program" the previous year in the *IEEE Transactions on Communications*. That provided the technical foundation of the "Internet," a term adopted later to refer to a network of networks that utilized ARPA's TCP/IP. In subsequent years the two men directed the development of Internet protocols in an environment they could control: the small community of ARPA contractors.

Cerf's departure marked a rift within the INWG. While Cerf and other ARPA con-

A BRIEF HISTORY OF THE OSI STANDARDS

1971: Cyclades packet-switching project launches in France.

1975: INWG submits a proposal to the International Telegraph and Telephone Consultative Committee (CCITT), which rejects it. Cerf resigns from INWG.

[BACHMAN] [ZIMMERMANN] [DAY]



1977: International Organization for Standardization (ISO) committee on Open Systems Interconnection is formed, with **Charles Bachman** as chairman.

January **1983:** U.S. Department of Defense's mandated use of TCP/IP on the ARPANET signals the "birth of the Internet."

May **1983:** ISO publishes "ISO 7498: The Basic Reference Model for Open Systems Interconnection" as an international standard.

1985: U.S. National Research Council recommends that the Department of Defense migrate gradually from TCP/IP to OSI.

1989: As OSI begins to founder, computer scientist Brian Carpenter gives a talk entitled "Is OSI Too Late?" He receives a standing ovation.

2013: IPv6 carries approximately 1 percent of global Internet traffic.

WWW

1991: Tim Berners-Lee announces public release of the WorldWideWeb application.

1974: IBM launches a packet-switching network called the Systems Network Architecture.



1974: Cerf and **Robert Kahn** publish "A Protocol for Packet Network Intercommunication," in *IEEE Transactions on Communications*.

1980: U.S. Department of Defense publishes "Standards for the Internet Protocol and Transmission Control Protocol."

1976: CCITT publishes Recommendation X.25, a standard for packet switching that uses "virtual circuits."

1988: U.S. market revenues for computer communications: \$4.9 billion.



1988: U.S. Department of Commerce mandates that government agencies buy OSI-compliant products.



1992: U.S. National Science Foundation revises policies to allow commercial traffic over the Internet.

1992: In a "palace revolt," Internet engineers reject the ISO ConnectionLess Network Protocol as a replacement for IP version 4.

IPv6

1996: Internet community defines IP version 6.

tractors eventually formed the core of the Internet community in the 1980s, many of the remaining veterans of INWG regrouped and joined the international alliance taking shape under the banner of OSI. The two camps became bitter rivals.

OSI was devised by committee, but that fact alone wasn't enough to doom the project—after all, plenty of successful standards start out that way. Still, it is worth noting for what came later.

In 1977, representatives from the British computer industry proposed the creation of a new standards committee devoted to packet-switching networks within the International Organization for Standardization (ISO), an independent nongovernmental association created after World War II. Unlike the CCITT, ISO wasn't specifically concerned with telecommunications—the wide-ranging topics of its technical committees included

TC1 for standards on screw threads and TC17 for steel. Also unlike the CCITT, ISO already had committees for computer standards and seemed far more likely to be receptive to connectionless datagrams.

The British proposal, which had the support of U.S. and French representatives, called for "network standards needed for open working." These standards would, the British argued, provide an alternative to traditional computing's "self-contained, 'closed' systems," which were designed with "little regard for the possibility of their interworking with each other." The concept of open working was as much strategic as it was technical, signaling their desire to enable competition with the big incumbents—namely, IBM and the telecom monopolies.

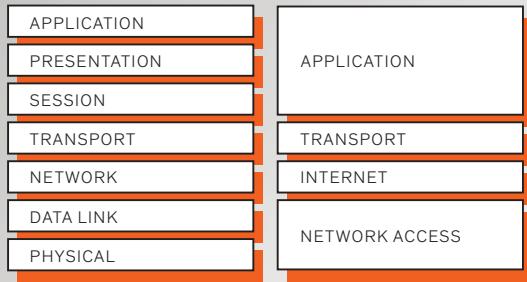
As expected, ISO approved the British request and named the U.S. database expert Charles Bachman as committee chairman. Widely respected in computer circles,

BARAN: COMPUTER HISTORY MUSEUM; DAVIES: NPL; MAP AND BACHMAN: COMPUTER HISTORY MUSEUM; ZIMMERMANN AND DAY: JOHN DAY; POUZIN: MARC WEBER/COMPUTER HISTORY MUSEUM; CERF: JOSE MERCADO/STANFORD NEWS SERVICE; MCKENZIE: ALEX MCKENZIE; KAHN: LOUIS F. BACHRACH

Bachman had four years earlier received the prestigious Turing Award for his work on a database management system called the Integrated Data Store.

When I interviewed Bachman in 2011, he described the "architectural vision" that he brought to OSI, a vision that was inspired by his work with databases generally and by IBM's Systems Network Architecture in particular. He began by specifying a reference model that divided the various tasks of computer communication into distinct layers. For example, physical media (such as copper cables) fit into layer 1; transport protocols for moving data fit into layer 4; and applications (such as e-mail and file transfer) fit into layer 7. Once a layered architecture was established, specific protocols would then be developed.

OSI vs. TCP/IP



A LAYERED APPROACH: The OSI reference model [left column] divides computer communications into seven distinct layers, from physical media in layer 1 to applications in layer 7. Though less rigid, the TCP/IP approach to networking can also be construed in layers, as shown on the right.

Bachman’s design departed from IBM’s Systems Network Architecture in a significant way: Where IBM specified a terminal-to-computer architecture, Bachman would connect computers to one another, as peers. That made it extremely attractive to companies like General Motors, a leading proponent of OSI in the 1980s. GM had dozens of plants and hundreds of suppliers, using a mix of largely incompatible hardware and software. Bachman’s scheme would allow “interworking” between different types of proprietary computers and networks—so long as they followed OSI’s standard protocols.

The layered OSI reference model also provided an important organizational feature: modularity. That is, the layering allowed committees to subdivide the work. Indeed, Bachman’s reference model was just a starting point. To become an international standard, each proposal would have to complete a four-step process, starting with a working draft, then a draft proposed international standard, then a draft international standard, and finally an international standard. Building consensus around the OSI reference model and associated standards required an extraordinary number of plenary and committee meetings.

OSI’s first plenary meeting lasted three days, from 28 February through 2 March 1978. Dozens of delegates from 10 countries participated, as well as observers from four international organizations. Everyone who attend-

ed had market interests to protect and pet projects to advance. Delegates from the same country often had divergent agendas. Many attendees were veterans of INWG who retained a wary optimism that the future of data networking could be wrested from the hands of IBM and the telecom monopolies, which had clear intentions of dominating this emerging market.

Meanwhile, IBM representatives, led by the company’s capable director of standards, Joseph De Blasi, masterfully steered the discussion, keeping OSI’s development in line with IBM’s own business interests. Computer scientist John Day, who designed protocols for the ARPANET, was a key member of the U.S. delegation. In his 2008 book *Patterns in Network Architecture* (Prentice Hall), Day recalled that IBM representatives expertly intervened in disputes between delegates “fighting over who would get a piece of the pie.... IBM played them like a violin. It was truly magical to watch.”

Despite such stalling tactics, Bachman’s leadership propelled OSI along the precarious path from vision to reality. Bachman and Hubert Zimmermann (a veteran of Cyclades and INWG) forged an alliance with the telecom engineers in CCITT. But the partnership struggled to overcome the fundamental incompatibility between their respective

worldviews. Zimmermann and his computing colleagues, inspired by Pouzin’s datagram design, championed “connectionless” protocols, while the telecom professionals persisted with their virtual circuits. Instead of resolving the dispute, they agreed to include options for both designs within OSI, thus increasing its size and complexity.

This uneasy alliance of computer and telecom engineers published the OSI reference model as an international standard in 1984. Individual OSI standards for transport protocols, electronic mail, electronic directories, network management, and many other functions soon followed. OSI began to accumulate the trappings of inevitability. Leading computer companies such as Digital Equipment Corp., Honeywell, and IBM were by then heavily invested in OSI, as was the European Economic Community and national governments throughout Europe, North America, and Asia.

Even the U.S. government—the main sponsor of the Internet protocols, which were incompatible with OSI—jumped on the OSI bandwagon. The Defense Department officially embraced the conclusions of a 1985 National Research Council recommendation to transition away from TCP/IP and toward OSI. Meanwhile, the Department of Commerce issued a mandate in 1988 that the OSI standard be used in all computers purchased by U.S. government agencies after August 1990.

While such edicts may sound like the work of overreaching bureaucrats, remember that throughout the 1980s, the Internet was still a *research* network: It was growing rapidly, to be sure, but its managers did not allow commercial traffic or for-profit service providers on the government-subsidized backbone until 1992. For businesses and other large entities that wanted to exchange data between different kinds of computers or different types of networks, OSI was the only game in town.

That was not the end of the story, of course. By the late 1980s, frustration with OSI’s slow development had reached a boil-



WHAT’S IN A NAME: At a July 1986 meeting in Newport, R.I., representatives from France, Germany, the United Kingdom, and the United States considered how the OSI reference model would handle the crucial functions of naming and addressing on the network.

ing point. At a 1989 meeting in Europe, the OSI advocate Brian Carpenter gave a talk titled “Is OSI Too Late?” It was, he recalled in a recent memoir, “the only time in my life” that he “got a standing ovation in a technical conference.” Two years later, the French networking expert and former INWG member Pouzin, in an essay titled “Ten Years of OSI—Maturity or Infancy?,” summed up the growing uncertainty: “Government and corporate policies never fail to recommend OSI as the solution. But, it is easier and quicker to implement homogenous networks based on proprietary architectures, or else to interconnect heterogeneous systems with TCP-based products.” Even for OSI’s champions, the Internet was looking increasingly attractive.

That sense of doom deepened, progress stalled, and in the mid-1990s, OSI’s beautiful dream finally ended. The effort’s fatal flaw, ironically, grew from its commitment to openness. The formal rules for international standardization gave any interested party the right to participate in the design process, thereby inviting structural tensions, incompatible visions, and disruptive tactics.

OSI’s first chairman, Bachman, had anticipated such problems from the start. In a conference talk in 1978, he worried about OSI’s chances of success: “The organizational problem alone is incredible. The technical problem is bigger than any one previously faced in information systems. And the political problems will challenge the most astute statesmen. Can you imagine trying to get the representatives from ten major and competing computer corporations, and ten telephone companies and PTTs [state-owned telecom monopolies], and the technical experts from ten different nations to come to any agreement within the foreseeable future?”

Despite Bachman’s and others’ best efforts, the burden of organizational overhead never lifted. Hundreds of engineers attended the meetings of OSI’s various committees and working groups, and the bureaucratic procedures used to structure the discussions didn’t allow for the speedy production of standards. Everything was up for debate—even trivial nuances of language, like the difference between “you will comply” and “you should comply,” triggered complaints. More significant rifts continued between OSI’s computer and telecom experts, whose technical and business plans remained at

odds. And so openness and modularity—the key principles for coordinating the project—ended up killing OSI.

Meanwhile, the Internet flourished. With ample funding from the U.S. government, Cerf, Kahn, and their colleagues were shielded from the forces of international politics and economics. ARPA and the Defense Communications Agency accelerated the Internet’s adoption in the early 1980s, when they subsidized researchers to implement Internet protocols in popular operating systems, such as the modification of Unix by the University of California, Berkeley. Then, on 1 January 1983, ARPA stopped supporting the ARPANET host protocol, thus forcing its contractors to adopt TCP/IP if they wanted to stay connected; that date became known as the “birth of the Internet.”

And so, while many users still expected OSI to become the future solution to global network interconnection, growing numbers began using TCP/IP to meet the practical near-term pressures for interoperability.

Engineers who joined the Internet community in the 1980s frequently misconstrued OSI, lampooning it as a misguided monstrosity created by clueless European bureaucrats. Internet engineer Marshall Rose wrote in his 1990 textbook that the “Internet community tries its very best to ignore the OSI community. By and large, OSI technology is *ugly* in comparison to Internet technology.”

Unfortunately, the Internet community’s bias also led it to reject any technical insights from OSI. The classic example was the “palace revolt” of 1992. Though not nearly as formal as the bureaucracy that devised OSI, the Internet had its Internet Activities Board and the Internet Engineering Task Force, responsible for shepherding the development of its standards. Such work went on at a July 1992 meeting in Cambridge, Mass. Several leaders, pressed to revise routing and addressing limitations that had not been anticipated when TCP and IP were designed, recommended that the community consider—if not adopt—some technical protocols developed within OSI. The hundreds of Internet engineers in attendance howled in protest and then sacked their leaders for their heresy.

Although Cerf and Kahn did not design TCP/IP for business use, decades of govern-

ment subsidies for their research eventually created a distinct commercial advantage: Internet protocols could be implemented for free. (To use OSI standards, companies that made and sold networking equipment had to purchase paper copies from the standards group ISO, one copy at a time.) Marc Levilion, an engineer for IBM France, told me in a 2012 interview about the computer industry’s shift away from OSI and toward TCP/IP: “On one side you have something that’s free, available, you just have to load it. And on the other side, you have something which is much more architected, much more complete, much more elaborate, but it is expensive. If you are a director of computation in a company, what do you choose?”

By the mid-1990s, the Internet had become the de facto standard for global computer networking. Cruelly for OSI’s creators, Internet advocates seized the mantle of “openness” and claimed it as their own. Today, they routinely campaign to preserve the “open Internet” from authoritarian governments, regulators, and would-be monopolists.

In light of the success of the nimble Internet, OSI is often portrayed as a cautionary tale of overbureaucratized “anticipatory standardization” in an immature and volatile market. This emphasis on its failings, however, misses OSI’s many successes: It focused attention on cutting-edge technological questions, and it became a source of learning by doing—including some hard knocks—for a generation of network engineers, who went on to create new companies, advise governments, and teach in universities around the world.

Beyond these simplistic declarations of “success” and “failure,” OSI’s history holds important lessons that engineers, policymakers, and Internet users should get to know better. Perhaps the most important lesson is that “openness” is full of contradictions. OSI brought to light the deep incompatibility between idealistic visions of openness and the political and economic realities of the international networking industry. And OSI eventually collapsed because it could not reconcile the divergent desires of all the interested parties. What then does this mean for the continued viability of the open Internet? ■

POST YOUR COMMENTS online at
<http://spectrum.ieee.org/osihistory0813>