

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/234792183>

The Past and Future History of the Internet

Article in Communications of the ACM · February 1997

Impact Factor: 3.62 · DOI: 10.1145/253671.253741

CITATIONS

125

READS

44

9 authors, including:



Vinton G. Cerf

Google Inc.

121 PUBLICATIONS 2,872 CITATIONS

SEE PROFILE



L. Kleinrock

University of California, Los Angeles

228 PUBLICATIONS 19,484 CITATIONS

SEE PROFILE



Lawrence Roberts

31 PUBLICATIONS 2,530 CITATIONS

SEE PROFILE



Stephen Wolff

Internet2

20 PUBLICATIONS 367 CITATIONS

SEE PROFILE

Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn,
Leonard Kleinrock, Daniel C. Lynch, Jon Postel,
Lawrence G. Roberts, Stephen S. Wolff

The Past and Future History of the **INTERNET**

the science of future technology

THE INTERNET HAS REVOLUTIONIZED THE COMPUTER AND COMMUNICATIONS world like nothing before. The telegraph, telephone, radio, and computer have all set the stage for the Internet's unprecedented integration of capabilities. The Internet is at once a worldwide broadcasting capability, a mechanism for information dissemination, and a medium for collaboration and interaction between individuals and their computers without regard for geographic location.

The Internet also represents one of the most successful examples of sustained investment and commitment to research and development in information infrastructure. Beginning with early research in packet switching, the government, industry, and academia have been partners in evolving and deploying this exciting new technology. Today, terms like "leiner@mcc.com" and "http://www.acm.org" trip lightly off the tongue of random people on the street.¹

The Internet today is a widespread information infrastructure, the initial prototype of what is often called the National (or Global or Galactic) Information Infrastructure. Its history is complex and involves

many aspects—technological, organizational, and community. And its influence reaches not only to the technical fields of computer communications but throughout society as we move toward increasing use of online tools to accomplish electronic commerce, information acquisition, and community operations.²

Origins

The first recorded description of the social interactions that could be enabled through networking was a series of memos written August 1962 by J.C.R. Licklider of MIT, discussing his "Galactic Network" concept [6]. Licklider envisioned a globally interconnected set of

¹Perhaps this is an exaggeration due to the lead author's residence in Silicon Valley.

²For a more detailed version of this article, see <http://www.isoc.org/internet-history>.

computers through which everyone could quickly access data and programs from any site. In spirit, the concept was much like the Internet today. While at DARPA,³ he convinced the people who would be his successors there—Ivan Sutherland, Bob Taylor, and MIT researcher Lawrence G. Roberts—of the importance of this networking concept.

Leonard Kleinrock of MIT published the first paper on packet switching theory in July 1961 [5]. Kleinrock convinced Roberts of the theoretical feasibility of communications using packets rather than circuits—a major step toward computer networking. The other key step was to make the computers talk to each other. Exploring this idea in 1965 while working with Thomas Merrill, Roberts connected the TX-2 computer in Massachusetts to the Q-32 in California through a low-speed dial-up telephone line [8], creating the first-ever (though small) wide-area computer network. The result of this experiment: confirmation that time-sharing computers could work well together, running programs and retrieving data as necessary on remote machines, but that the circuit-switched telephone system was totally inadequate for the job. Thus confirmed was Kleinrock's conviction of the need for packet switching.

In late 1966, Roberts went to DARPA to develop the computer network concept and quickly put together a plan for the ARPANET, publishing it in 1967 [7]. Bolt, Beranek and Newman Corp. (BBN), under Frank Heart's leadership, developed the ARPANET switches (called IMPs), with Robert Kahn responsible for overall system design. Howard Frank and his team at Network Analysis Corp. worked with Roberts to optimize the network topology and economics. Due to Kleinrock's early development of packet switching theory and his focus on analysis, design, and measurement, his Network Measurement Center at UCLA was selected as the first node on the ARPANET. All this came together September 1969 when BBN installed the first switch at UCLA and the first host computer was connected. In December 1970, the Network Working Group (NWG) under Steve Crocker finished the initial ARPANET host-to-host protocol, called the Network Control Protocol (NCP). As the ARPANET sites completed implementing NCP during 1971–1972, network users finally could begin to develop applications.

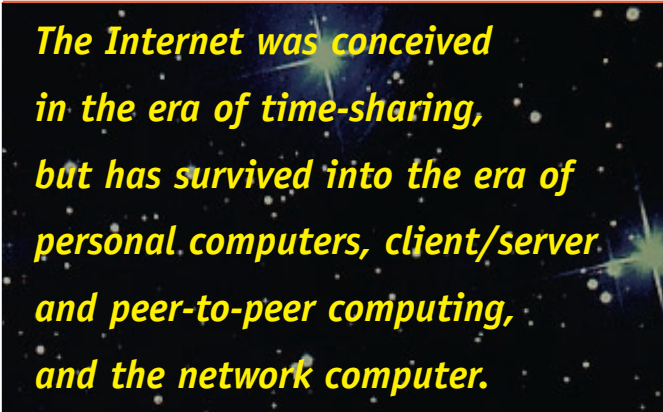
In October 1972, a large, successful demonstration of the ARPANET took place—the first public demon-

stration of this new network technology. Also in 1972, electronic mail, the initial “hot” application, was introduced. In March, Ray Tomlinson of BBN wrote the basic email message send-and-read software, motivated by ARPANET developers' need for an easy coordination mechanism. From there, email took off as the most popular network application and as a harbinger of the kind of people-to-people communication activity we see on the World-Wide Web today.

Initial Internetting Concepts

The original ARPANET grew into the Internet based on the idea that there would be multiple independent networks of rather arbitrary design. Beginning with the ARPANET as the pioneering packet-switching network, it soon grew to include packet satellite networks, ground-based packet radio networks, and other networks. Today's Internet embodies a key underlying technical idea: open-architecture networking. In this approach, the choice of any individual network technology is not dictated by a particular network architecture but can be selected freely by a provider and made to interwork with the other networks through a meta-level “internetworking architecture.” Each network can be designed to fit a specific environment and user requirements.

The idea of open-architecture networking—introduced by Kahn in late 1972 shortly after arriving at DARPA—was guided by four critical ground rules:



*The Internet was conceived
in the era of time-sharing,
but has survived into the era of
personal computers, client/server
and peer-to-peer computing,
and the network computer.*

- Each distinct network had to stand on its own, and no internal changes could be required of any such network before being connected to the Internet.
- Communications would be on a best-effort basis. If a packet didn't make it to the final destination, it would quickly be retransmitted from the source.
- Black boxes (later called gateways and routers)

³The Advanced Research Projects Agency (ARPA) changed its name to Defense

- mation would be retained by the gateways about individual flows of packets passing through them, keeping them simple and avoiding complicated adaptation and recovery from various failure modes.
- There would be no global control at the operations level.

Kahn first began work on a communications-oriented set of operating system principles while at BBN [4]. After joining DARPA and initiating the Internet program, he asked Vinton Cerf (then at Stanford University) to work with him on the detailed design of the protocol. Cerf had been deeply involved in the original NCP design and development and was already knowledgeable about interfacing to existing operating systems. So, armed with Kahn's architectural approach to communications and with Cerf's NCP experience, these two teamed up to spell out the details of what became the Transmission Control Protocol/Internet

The most pressing question for the future of the Internet is not how the technology will change, but how the process of change and evolution itself will be managed.

Protocol (TCP/IP).

The original Cerf/Kahn paper [1] on the Internet described a protocol, called TCP, providing all the Internet's transport and forwarding services. Kahn had intended that TCP would support a range of transport services—from the totally reliable sequenced delivery of data (virtual circuit model) to a datagram service in which the application made direct use of the underlying network service, a process that could imply occasional lost, corrupted, or reordered packets.

However, the initial effort to implement TCP resulted in a version allowing only virtual circuits. This model worked fine for file transfer and remote login applications, but some of the early work on advanced network applications, particularly packet voice in the 1970s, made clear that in some cases packet losses should not be corrected by TCP but left to the application to deal with. This insight led to a reorganization of the original TCP into two proto-

forwarding of individual packets and the separate TCP concerned with such service features as flow control and recovery from lost packets. For applications that did not want the services of TCP, an alternative called the User Datagram Protocol (UDP) was added to provide direct access to the basic IP service.

In addition to email, file transfer, and remote login, other applications were proposed in the early days of the Internet, including packet-based voice communication (the precursor of Internet telephony), various models of file and disk sharing, and early "worm" programs illustrating the concept of agents (and viruses). The Internet was not designed for just one application but as a general infrastructure on which new applications could be conceived, exemplified later by the emergence of the Web. The general-purpose nature of the service provided by TCP and IP makes this possible.

Proving the Ideas

DARPA funded three efforts to implement TCP: Stanford (Cerf), BBN (Tomlinson), and University College London (Peter Kirstein). The Stanford team produced a detailed specification, yielding within about a year three independent interoperable implementations of TCP.

This was the beginning of long-term experimentation and development of Internet concepts and technology—along with the constituent networking technologies [3]. Each expansion has involved new challenges. For example, the early implementations of TCP were done for large time-sharing systems. When desktop computers first appeared, it was thought by some that TCP was too big and complex to run on a personal computer. But David Clark and his research group at MIT produced an implementation first for the Xerox Alto (the early personal workstation developed at Xerox PARC) and then for the IBM PC, showing that workstations, as well as large time-sharing systems, could be part of the Internet.

Widespread development of local-area networks (LANs), PCs, and workstations in the 1980s allowed the nascent Internet to flourish. Ethernet technology (developed by Bob Metcalfe at Xerox PARC in 1973) is now probably the dominant network technology in the Internet and PCs and workstations the dominant computers. The increasing scale of the Internet also resulted in several new approaches. For example, the Domain Name System was invented (by Paul Mockapetris, then at USC's Information Sciences Institute) to provide a scalable mechanism for resolving hierarchical host names (e.g., www.acm.org) into Internet

approaches led to a hierarchical model of routing, with an Interior Gateway Protocol (IGP) used inside each region of the Internet and an Exterior Gateway Protocol (EGP) used to tie the regions together. New approaches for address aggregation, particularly classless interdomain routing (CIDR), were recently introduced to control the size of router tables.

Another major challenge was how to propagate the changes to the software, particularly host software. DARPA supported the University of California at Berkeley to investigate modifications to the Unix operating system, including incorporating TCP/IP developed at BBN. Although Berkeley later rewrote the BBN code to more efficiently fit into the Unix system and kernel, incorporation of TCP/IP into the Unix BSD system proved a critical element in dispersing the protocols to the research community. Much of the computer science community began using Unix BSD in their day-to-day computing environments. Looking back, the strategy of incorporating Internet protocols into a supported operating system for the research community was a key element in the Internet's successful widespread adoption.

TCP/IP was adopted as a defense standard in 1980, enabling the defense community to begin sharing the DARPA Internet technology base and leading directly to the partitioning of the military and non-military communities. By 1983, ARPANET was being used by a significant number of defense R&D and operational organizations. The transition of ARPANET from NCP to TCP/IP in 1983 permitted it to be split into a MILNET supporting operational requirements and an ARPANET supporting research needs.

Thus, by 1985, the Internet was established as a technology supporting a broad community of researchers and developers and was beginning to be used by other communities for daily computer communications. Email was being used broadly across several communities, often with different systems, demonstrating the utility of broad-based electronic communications between people.

Transition to Widespread Infrastructure

At the same time Internet technology was being experimentally validated and widely used by a subset of computer science researchers, other networks and networking technologies were being pursued. The usefulness of computer networking—especially email—demonstrated by DARPA and Department of Defense contractors on the ARPANET was not lost on

1970s computer networks began springing up wherever funding was found for the purpose, including the Department of Energy's MFENET and HEPNET, NASA's SPAN, the computer science community's CSNET, the academic community's BITNET, and USENET based on Unix UUCP protocols. Commercial networking technologies were being pursued as well, including IBM's SNA, Xerox's XNS, and Digital Equipment Corp.'s DECNET.

It remained for the British JANET (1984) and U.S. NSFNET (1985) programs to explicitly announce their intent to serve the entire higher education community, regardless of discipline. In 1985, Dennis Jennings came from Ireland for a year to lead the National Science Foundation's NSFNET program. He helped NSF make a critical decision—that TCP/IP would be mandatory for NSFNET. And when Stephen Wolff took over the NSFNET program in 1986, he recognized the need for a wide-area networking infrastructure to support the general academic and research community, as well as the need to develop a strategy for establishing such infrastructure to ultimately be independent of direct federal funding.

Policies and strategies were adopted to achieve that end. So while federal agencies shared the cost of common infrastructure, such as trans-oceanic circuits, NSF encouraged regional networks of the NSFNET to seek commercial, non-academic customers. And NSF enforced an acceptable-use policy, prohibiting Backbone use for purposes “not in support of research and education.” The predictable (and intended) result of encouraging commercial network traffic at the local and regional levels, while denying its access to national-scale transport, was the emergence and growth of “private,” competitive, long-haul networks, such as PSI, UUNET, ANS CO+RE, and (later) others.

NSF's privatization policy culminated in April 1995 with the defunding of the NSFNET Backbone. The funds thereby recovered were (competitively) redistributed to regional networks to buy national-scale Internet connectivity from the now numerous, private, long-haul networks. The Backbone had made the transition from a network built from routers out of the research community (David Mills's “Fuzzball” routers) to commercial equipment. In its eight-and-a-half-year lifespan, the Backbone had grown from six nodes with 56Kbps links to 21 nodes with multiple 45Mbps. It also saw the Internet grow to more than 50,000 networks on all seven continents and outer space (with 29,000 networks in the U.S.).

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.