# Who Really Invented the Internet?

**Leo L. Beranek**, Cambridge, Massachusetts

**Before Ebay, before Google, before Al Gore, before Tim Berners-Lee – there was a small acoustical consulting firm in Cambridge, MA that put together the team: that designed the system, that fabricated the hardware, that wrote the code, that built the house of Internet.**

On October 3, 1969, for the first time, two computers at remote locations 'spoke' to each other over the roadbed of the Internet. Connected by 350 miles of leased telephone line, the two machines, one at the University of California in Los Angeles and the other at Stanford Research Institute, attempted the simplest of messages: the word 'login' transmitted one letter at a time. 'L' and 'O' transmitted perfectly. When the 'G' was transmitted, the SRI computer crashed. Despite the crash, a major hurdle had been cleared and the computers had actually managed to convey a meaningful message, even if not the one planned; in its own phonetic fashion, the UCLA computer said 'ello' to its compatriot in Stanford. The first, albeit tiny, computer, innovative network was now in operation.

With very little danger of negation, one can assume that the Internet is one of the five prodigious inventions of the twentieth century, rubbing shoulders with television, aircraft, atomic energy and space exploration. Unlike several of those, however, it did not have its beginnings in the nineteenth century. As late as 1940, not even an imagination like that of Jules Verne could have foreseen how the collaboration of physical scientists and psychologists in World War II would culminate three decades later in the beginnings of a new communication revolution. Even the blue-ribbon laboratories of A.T.&T., IBM, GE, and the like, when presented with the prospect of a group of computers that could speak simultaneously over a maze of wires, could imagine nothing more than a mechanism that depended on computer-to-computer communication over a single telephone line using central-office switching methods. The vision that went further came from a few other institutions and companies and, most important, the individuals working at them

While one can view the October 1969 transmission as a beginning, for those researchers working in communications and artificial intelligence in the preceding decades, it was an event with long and complicated roots. This article will trace those beginnings from their origin in World War II voice-communication laboratories and seek to demonstrate how the conceptual leaps of a number of gifted individuals, as well as their hard work and production skills, made possible the e-mail you receive each day. As hard as it may be to pin down something as nebulous as invention, the first network is not hard to identify. The computer in Los Angeles said 'ello' to the computer at Stanford via a tiny packet-switched network called ARPANET, named for the Advanced Research Projects Agency in the U.S. Department of Defense. Bolt Beranek and Newman, ARPANET's creator and manager for 20 years, owed its success to several factors – proximity to two renowned universities, a dedication to hiring only the best minds and the free-spending-on-research policies of the US Government following the advent of Sputnik.

In 1948, Professors Richard Bolt, Robert Newman, and I, with the blessing of the Massachusetts Institute of Technology, formed the acoustical consulting firm, Bolt Beranek and Newman (BBN), then a partnership. Little did we know it at the time, but we were setting the foundation for the development of the Internet, the genesis of which required three conceptual innovations – man-machine systems or symbiosis, time sharing and packet-switching. Over the next fifteen years, BBN would bring together the minds capable of conceiving these three and making them work.

In retrospect, the most resonant of the three for the computer-literate non-specialist appears to be "man-machine symbiosis," a ground breaking concept articulated largely by J. C. R. Licklider. He envisioned access to large computers, then common in major universities, by nearby users who would employ those resources to solve every type of problem. His paper in 1960, while at BBN, was an important rung in the ladder that was to establish his name as the forefather of the Internet. His summary contains:

> *. . . Men will set the goals, formulate the hypotheses, determine the criteria and perform the evaluations. Computing machines will do the routine work that must be done to prepare the way for insights and decisions in technical and scientific thinking . . . Prerequisites for the achievement of an effective, cooperative association include developments in computer time sharing, in memory components, in memory organization, in programming languages and in input and output equipment.*

There were, however, twelve years between the founding of BBN and Licklider's insight. Furthermore, the forces that brought BBN and Licklider together reached back to World War II, when the Psycho-Acoustic Laboratory (PAL) and the Electro-Acoustic Laboratory (EAL) at Harvard tackled some problems that prevented communication in and from U.S. Air Force aircraft flying at high altitudes. The EAL, under my direction, collaborated closely with PAL, where a young scientist named Licklider demonstrated impressive proficiency in both physics and psychology. We solved the voice-communication problem by finding and developing microphones, earphones and amplifiers with altitude-friendly components and proving their effectiveness through psycho-acoustic tests. It is apparent now that the daily closely-wedded cooperation between a group of physicists and a group of psychologists that took place at these two separate laboratories was unique in history.

At the end of World War II, I migrated to MIT, taking on the title Associate Professor of Communication Engineering and Technical Director of the Acoustics Laboratory. With top-flight post-doctorates hireable, Bolt Beranek and Newman soon expanded its efforts to include research in psychoacoustics. Research support for them was obtained from government agencies for: (1) speech compression; (2) criteria for prediction of speech intelligibility in noise; (3) the effects of noise on sleep; and (4) life sciences. In 1949, I influenced the administration of MIT's Department of Electrical Engineering to appoint Licklider as an Associate Professor to work with me in the Acoustics Laboratory on voice communication problems.

In about 1956, I decided that BBN ought to expand its efforts to include "man-machine systems," envisioning work on aircraft instrument landing systems and machines that would efficiently amplify human labor. With this goal, I again found it necessary to attract new talent. This time, I wanted an outstanding experimental psychologist to head up the activity, preferably one acquainted with the then rudimentary field of digital computers. It is not surprising that J. C. R. Licklider became my top candidate for heading such a new division. Of course, a position at BBN would mean that he must give up a tenured faculty position. One of the strengths of my discussion was an offer of stock options with the hopeful promise that with time they could become an important source of wealth for him, which they did. Licklider came aboard BBN in the Spring of 1957 as a Vice President.

Licklider had been on board only a few months when he came to me and said that he wanted BBN to buy a computer for him and his group. I answered that we had a punched-card computer in the financial department and analog computers in the experimental psychology group. He said that they were of no interest to him, he wanted a digital computer that used punched paper tapes for its external memories. 'What do you have in mind?' I asked. He said that he had been negotiating with the Royal-McBee company and they would give us a discounted price on their latest computer. "What will it cost?" was my next question. He said, 'Around $30,000." I said that we had never spent anything approaching that amount of money on a single research apparatus. "What are you going to do with it?" I queried. "I don't know, but if BBN is going to be an important company in the future, it must be in computers."

Lick almost floored me with the $30,000 answer and no plan for

*J. C. R. Licklider, psychophysicist who advocated time-shared computers and transmission lines. (1965) – Courtesy of Koby, Cambridge, MA.*

using it. But Lick was persistent. He believed we would be able to obtain contracts from the government that would make use of a late model digital computer, and we did. I concluded that if Lick had strong convictions that a computer was the right way to go, and with his already national reputation for doing good research, it was a risk worth taking. My partners were apprehensive since they did not know Licklider well. But, we bit the bullet and Lick received permission to make the purchase. We also encouraged him to hire suitable staff.

The new computer was a vacuum-tube machine with no core or vacuum-tube memory, only a drum. It used punched paper tape for longer memory, but paper tape was unreliable. It made mistakes and was very slow. But, Lick, with his Mona-Lisa grin, said that it was fun to play with. For awhile he devoted almost all his time to learning how to write programs for a digital device and deciding on what research could be performed with it.

Less than a year after that computer's arrival, Kenneth Olsen, the president of the fledgling Digital Equipment Corporation (DEC) stopped by ostensibly to see our computer. After chatting with us awhile and satisfying himself that Lick really understood digital computation, he asked if we would consider a project. He explained that DEC had just completed construction of a prototype digital computer, which DEC named PDP-1. It was based on a new design made at MIT's Lincoln Laboratory. He asked us to become a test site for a month. We consented.

**The new computer arrived in 1960, with a price tag of nearly $150,000.**

Shortly thereafter the prototype PDP-1 arrived. It was enormous, compared to the Royal-McBee, and we had no place to put it except in the main visitors' lobby. To hide it from strangers we put Japanese screens around it. Lick and Ed Fredkin, principally, and several others, put it through its paces for most of a month. At the end, Lick gave Olsen a list of suggestions including how to make it more user-friendly to researchers – thin skinned, as he called it, so that researchers could get inside and make changes.

We liked the computer and made arrangements with Olsen to sell us their first PDP-1, on a standard lease basis. The new computer arrived in 1960, with a price tag of nearly $150,000, and we immediately put it to work on government-supported research projects. Our personnel in computer research and systems, most hired before 1969, included many names, some now famous in the field. We began to be known in Cambridge as "The Third University." Some researchers thought BBN was a better place to work than a university because there was no teaching and no committee work.

Between 1960 and 1962, with his new PDP-1 on board and several more coming in, Lick turned his attention to time sharing. Time sharing, means simultaneous use of a machine by two or more persons; it works in such a way that one operator need never know that others were also busily using the same computer. It is in a sense, the precursor of the in-house servers so many businesses and institutions depend on today. Lick alone, however, was not responsible for the development of time sharing work at BBN. Always an astute manager, in 1962 he brought in John McCarthy, whom he knew from his days at MIT, and Marvin Minsky of Harvard. McCarthy and Minsky were primarily interested in research on artificial intelligence, and we had a contract on that subject. Lick engaged them as consultants, nearly full-time during the summer of 1962.

McCarthy and Minsky quickly produced results at BBN, though McCarthy has made a point of further sharing the credit, as he explained in 1989: "I was promoting time sharing in general. The person who said that time sharing could be done on a small computer, namely a PDP-1, was Ed Fredkin. I kept arguing with him. I said that an interrupt system was needed. And he said, "We can do that." Also needed was some kind of swapper. "We can do that." . . . Basically, Fredkin had done all the work and then I supervised Sheldon Boilen in actually implementing that time sharing system."

When the time-sharing program was finished, the computer screen on the PDP-1 was divided into four parts, each part assigned to a user. A public demonstration of time sharing, the first ever, was made by BBN in the fall of 1962 with one operator in Washington, D.C. and two in Cambridge, MA.

Concrete applications of this major advance followed rapidly. BBN in 1962 undertook to design, develop and implement a time-shared information system for the Massachusetts General Hospital with which nurses and doctors could keep hospital records in a central computer. BBN formed a subsidiary company, TELCOMP, to perform a time sharing service in the greater Boston and New York areas with paying subscribers who used Teletypewriters connected to dial-up telephone lines to access our time-shared DEC PDP computers.

With the PDP-1 onboard, BBN was able to hire first-class people from MIT and Harvard. By 1961 we had leased two more PDP-1s and a fourth, specially built for time sharing, arrived in 1963. As the decade wore on we purchased advanced computers from DEC, IBM and SDS (Scientific Data Systems), and also separate large-disk memories. By 1968 BBN had over 600 employees, over half in the computer division. BBN was said at that time to have more prime contracts with Federal agencies than any other company in Massachusetts.

In October 1962, Licklider was lured away from BBN for a one-year stint, which stretched to two, by the Advanced Research Projects Agency of the Department of Defense, ARPA. Within this office, Lick took the title, "Director of Behavioral Sciences." By the end of Lick's two-year stint, ARPA had spread time sharing nationwide through contract awards.

## ARPANET

After Lick's term at ARPA, the directorship eventually passed to Robert Taylor who served from 1966 to 1968 and who oversaw the beginnings of ARPA's plan to build a packet-switched network. In 1967, he was largely responsible for recruiting Lawrence Roberts from MIT's Lincoln Laboratory as a Program Manager with responsibility for the network project. The real challenge of Robert's job was developing a system for connecting a matrix of computers with different characteristics. The stated purpose was to enable smaller research laboratories to access large-scale computers at large research centers, without the necessity of ARPA supplying every laboratory with a multimillion dollar machine. Roberts' primary influences, whose ideas I will trace in a moment, were Wes Clark, Paul Baran, Leonard Kleinrock, Robert Kahn and Donald Davies.

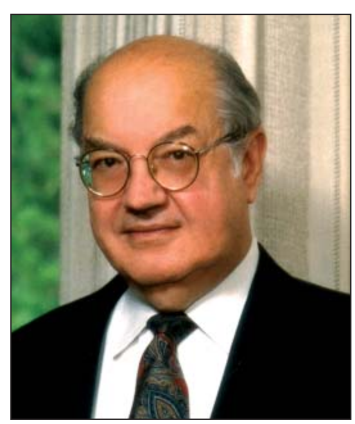

*Paul Baran. – Courtesy of Paul Baran.*

Roberts actually met with a general lack of enthusiasm when he initiated the discussion. Most said their computers were busy full time and they could think of nothing they would want to do cooperatively with other computer sites. But Roberts proceeded undaunted.

Wes Clark of Washington University in St. Louis proposed a different scheme than Roberts had first presented. He said the network should be composed of a set of identical mini-computers, which he called 'nodes,' all interconnected. All routing of message blocks would be in their charge. The owners of the large computers at the various locations, called 'hosts,' would need to construct ways to connect them to the network via these nodes. Roberts wrote a memorandum for circulation with Clark's network suggestion. He named the nodes, Interface Message Processors, IMPs. They were to perform the functions of interconnecting the network, breaking messages into packets, sending and receiving them, checking for errors, verifying that the packets arrive at their intended destinations, and reassembling the packets. Where did Roberts find the underlining information to write a meaningful request for bids from potential contractors? Paul Baran of the RAND Corporation, with US Air Force support, had tackled in 1960 the question of a vulnerable telephone communication system during a nuclear attack. He created the concept of dividing messages into "message blocks" that would be sent over a network of telephone lines. He visualized that successive blocks would originate in the computer 'nodes' and would travel over different routes (telephone lines) and be reassembled at their destination, another computer node. In other words, he had visualized a means for time-sharing the transmission lines by interleaving the message blocks originating from the various nodes, each of which were connected to other nodes by as many as eight telephone lines.

In RAND's attempt to sell the idea to A.T.&T., Baran answered their questions in a series of 11 volumes, compiled from 1960 to 1965. A.T.&T. continued to resist the idea, The struggle ended in the USAF files. Roberts discovered and digested those volumes. In 1968, Robert Taylor stated, "It was Paul, of course, who invented packet switching before there was any consideration of ARPANET."

Another contributor to the literature was Leonard Kleinrock, now at UCLA, who had written his MIT doctoral thesis in 1962 on communication networks. He had developed some design procedures, routing procedures and topology design. His work was known to Roberts as they had been students together at MIT, and was helpful after the ARPANET came into existence. Donald Davies of the National Physical Laboratory in the U.K., after observing MIT's Project MAC time-sharing system, proposed, in 1965, splitting typewritten messages into data 'packets,' each of the same length (about ten words) and time-sharing them on a single line, which he called "packet switching." He proved the elementary feasibility of his proposal by an experiment over one line in his laboratory. But nothing further came of this activity. Apparently, he was not aware of Baran's prior work.

By 1967, Roberts had become well acquainted with both Baran's and Davies' endeavors. Roberts chose Davies name 'packet' instead of Baran's "message block," probably because it is more economical when combined with 'switching.'

Bob Kahn at BBN had also played a part in Robert's early thinking about the network. During various early contacts, he and Kleinrock had convinced Roberts of the need for a full scale network using long distance telephone lines, rather than just a laboratory experiment, to show that packet switching was feasible.

Pulling together these insights, Roberts decided that ARPA should pursue what he dubbed "a packet switching network." The prospect would be a challenging one, largely because of the uncertainty of the overall design and the danger of failure. He said that old-line telephone engineers thought the idea unworkable. Some of the big companies maintained that the packets would circulate forever, making the whole effort a waste of time and money. Besides, they argued, why would anyone want such a network when the telephone system in the U.S. is the best in the world?

When Roberts released his "request for proposal" (RFP) in the summer of 1968, the door opened for BBN to step into the creation of the network. The request called for a trial network made up of four IMPs, connected to four host computers. If the four-node network proved itself, the network would expand to include fifteen more hosts. Although 140 companies received his request and thirteen submitted proposals, BBN won the contract. Related contracts for network "host sites" would go to UCLA, SRI, UCSB and the University of Utah. Even placing a bid for the contract for developing and overseeing the network, however, required a tremendous effort on the part of the BBN staff.

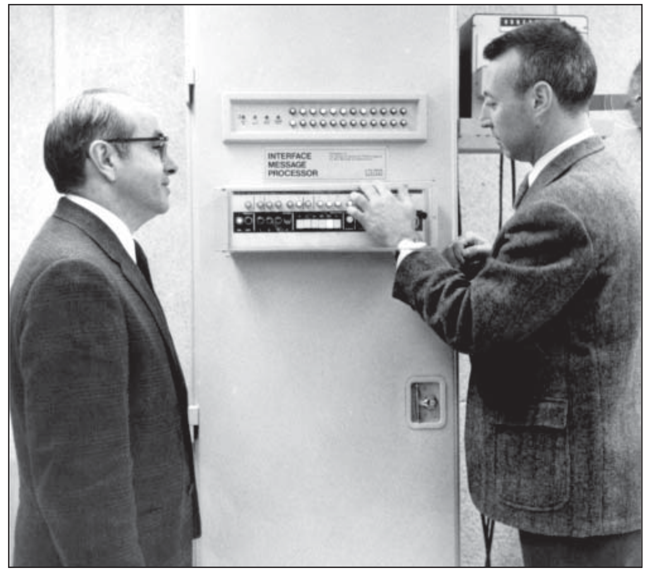


*Frank Heart (right) explaining the first Interface Message Processor (IMP) to BBN Vice President Samuel Labate, (1969) – Courtesy of Frank Heart.*

The "Request for Bid" arrived at BBN in August 1968. We in management picked Frank Heart to administer the bid proposal because he had already formed a team. The team had become experts at connecting electrical measuring devices through telephone lines for the gathering of information and thus they were pioneers in the construction of computing systems that worked in "real time" (versus. first recording data and later analyzing it). But the ARPANET was a new challenge, because it was filled with risk and with insufficient time for planning. This troubled him, because there was no certainty that a packet-switched network could ever be made to work. We in management thought that BBN should make every attempt to push ahead the boundaries of the unknown, and strongly urged response. Frank accepted the challenge. The time for response to the RFP was ridiculously short – nobody would plan on a decent night's sleep for the next month.

The ARPANET group worked until nearly dawn, day after day. An immediate task was the choice of a suitable computer for the IMPs. Heart insisted that reliability of operation had to be paramount in the design. He favored a newly-released computer, the Honeywell DDP-516, which was the size of a narrow refrigerator, encased in a thick-steel cabinet, and drop-tested for Navy use. Further, it was of the right digital capacity and was able to handle input and output signals with speed and efficiency. And very important, the manufacturing plant was a short drive away from BBN's offices. The proposal, the bid for the project, was finished on time. It filled two hundred pages, and appeared to demonstrate that BBN could surmount the problems inherent in the ARPA Request

In this proposal, the length of each packet was to be 1000 bits (approximately the equivalent of about 12 printed words). Every detail was written down for Larry Roberts to see, including how the network would address and queue the packets, determine the best transmission routes – all route-availability continually updated – handle congestion, recover from line, power and IMP failures, monitor and debug the machines from a remote-control center, and so on. They had discovered that the speed at which packets could be processed (end-to-end on the network) was about 1/10 of that called for in the RFP, i.e., ARPA's 1/2 second became 1/20th. Even so, BBN wrote in its proposal the hedging statement, "We take the position that it will be difficult to make the system work."

Thirteen companies submitted proposals, of which only two made the final list. BBN received on December 23, 1968 a telegram from Senator Ted Kennedy, ". . . *to be congratulated on winning the contract for the interfaith (sic) message processor.*" Roberts would remark that BBN had won the contract because of their well prepared proposal. He also believed the environment of a small company would be easier for him to work with.

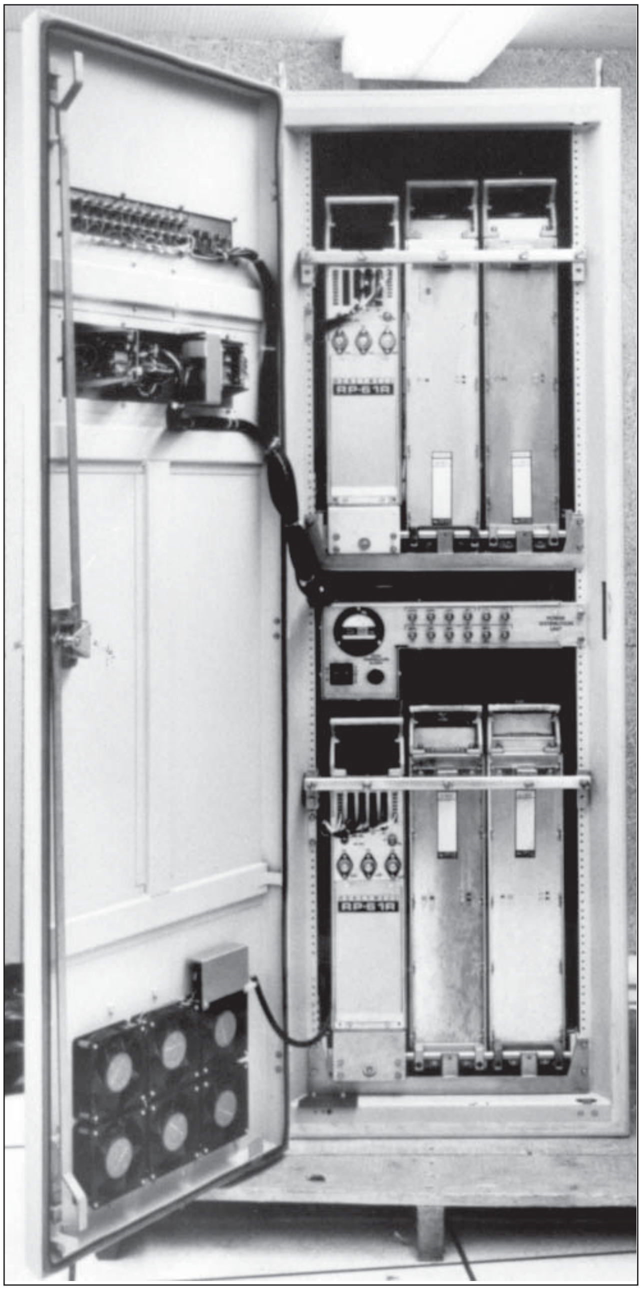

Interface Message Processor (IMP) with front door open. On the rear of the six inner units are thousands of wires. (1969) – Courtesy of Frank Heart.

As much work as the BBN staff had put into creating a bid, it was nothing compared to the work that came next – designing and building a new communication network, the basic elements of which were the Interface Message Processors that would be located at each host site. Although the goal at this point would only be the first IMP of a four-host demonstration network, the eight-month deadline still forced BBN into many weeks of late nights and marathon stints. Between New Year's Day and September 1, 1969, BBN would have to design the overall system and determine the network's hardware and software needs; acquire and modify the hardware; develop and document procedures for the host sites; ship a computer first to UCLA and the next three to the other venues within three months; and, finally, to oversee the arrival and behavior of each as it began service

The IMP team – so-called because the network was to be composed of 19 IMPs (to which host computers at as many establishments that were to be connected) – concentrated initially on obtaining a modified Honeywell computer. A first task was to design electronic attachments to the computer that would transfer electrical signals (electrical equivalents of the signals from a keyboard) into it and would receive from it responsive signals

*The creators of the IMP were (crouching, left to right) James Geisman, David Walden. Will Crowther; (next row) Truet Thach, William Bertell, Frank Heart, Marty Thorpe, Severo Ornstein, Robert Kahn; (rear) Ben Barker; (not pictured) Bernie Cosell. (1969) – Courtesy of Frank Heart.*

(like those that go out on today's Internet). These are equivalent to hearing and speech in human beings.

This computer had neither a hard drive, that is to say, a large writable memory, nor a floppy drive (such memory drives were not available). The total memory was only 12,000 bytes, compared to 100,000,000,000 bytes for today's PCs. The operating system (like Windows®) and programs were put into it from punched paper tapes.

Software writing in those days was impossibly difficult. It was very easy to get lost and concise programs were hard to write. The process was described thus, "The electrical signals from a standard Teletype machine were connected to a time-shared PDP-1, which created a computer 'code.' In the PDP-1, the code also was converted into the language of the Honeywell computer and transferred onto a paper tape by a punching device. The information on the tape, using the optical reader, would then be fed into the computer and the results would be observed on a screen."

### The Hardware

Each IMP as finally configured was responsible for lost or duplicate packets, and it checked regularly for telephone line failures, dead next-door IMPs, hosts that weren't operating and destinations that couldn't be reached because of line or IMP failures. It was also able to recover from its own failures, that is to say, it would restart immediately if the AC power failed. Even the memory would come through such a power line loss. Thus, BBN gave amazing 'survival' features to this distributed network.

The creation of a necessary 'checksum' also took place in the hardware. When a packet of information was transmitted from one IMP to the next, an error detecting number, the length of a packet, called the checksum, was transmitted with it. At the receiving point the checksum was recalculated by its IMP. If there was a difference, an error in transmission had occurred and the packet had to be resent. How to do this? The group decided to conserve on network space by having the receiving IMP acknowledge a correct transmission only if the checksum was correct. No acknowledgment meant that the packet had to be resent from the source IMP. How large should this checksum be? Analysis showed that a checksum equivalent to 24 bits (about three words) would be necessary for high accuracy, and this was adopted. This IMP-to-IMP error detection scheme was largely BBN's own invention, and was not performed in the software because its execution would have been too slow.

### The Software

The software group discovered how to make a very short and fast program to accept packets sent into the IMP from one line, determine where they should be sent, and send them out to the next IMP. The solution was unusually efficient for that time. The program used only 150 lines for the process of passing a packet through an IMP.

Next came the design of the routing system. Each packet had to be relayed from one IMP to another until it reached the IMP at its destination. This would have been simpler if it weren't for the need for a packet to find an alternate path [packet switching] in case the first path was busy or broke down. The group's final dynamic routing procedure was ingenious.

Other problems pertaining to the flow of the packets had to be solved. It was alright to say that the packets had to bounce across a sequence of IMPs to get to a destination, but what about traffic jams? One trouble was that after a packet or a series of packets had been sent, the next packet was not sent until acknowledgment of successful delivery of the first one(s) was received at the sending IMP. To explain – the total message might consist of many packets, and the incomplete set had to be held in a 'buffer' (storage space) at the receiving IMP until the whole message was assembled before it could be passed on to the host. Because the packets from different messages are inter-leaved, several different messages might be resting in partially finished state in any one IMP buffer. If a buffer became full, additional packets could not enter and no message could be completed. The solution – before the first packets in a message could be sent, the sending IMP would call ahead and reserve space for the complete message in the buffer of the receiving IMP. Other messages would have to be delayed until there was unreserved space. Fortunately, such delays are minuscule, because computers are speedy and transmission times are short.

## Report No. 1822

BBN's contract was only to develop the network to which host computers, other peoples' responsibility, would be connected. Nonetheless, it was apparent to the group that procedures would have to be spelled out for connection of the host computers to the IMPs. That is to say, a precise document would have to be written, so that those at the host locations could successfully construct electronics and software to go between their diverse makes of computers and the IMPs. Many discussions about such a document consumed precious time. It took two months to complete the document. The specification, known as BBN Report 1822, was completed during late spring. Kleinrock later articulated, "Anybody who was involved in the ARPANET will never forget that report number because it was the defining spec for how the things would mate."

The first big crisis came when the first Honeywell machine, Serial No. 0, was delivered to BBN, presumably modified to include all the features necessary to accommodate the software and perform the duties of an IMP on the network. But No. 0 didn't work. Debugging took months. The wiring to many of the pins had to be changed. A precise record had to be kept and Honeywell engineers informed of what modifications were necessary. The next machine from Honeywell would be Serial No. 1, and when received should be ready to be checked and shipped to UCLA.

The next crisis was approaching. Timing was now critical if the Labor Day deadline was to be met. Many weeks would pass between the time when BBN had communicated the necessary changes to Honeywell and when Serial No. 1 arrived, only two weeks before it would have to be shipped to California. But Serial No. 1 also did not work. Considerable rewiring was necessary. Finally, the completed machine worked properly, but occasionally it crashed, once a day or so when fully loaded. This behavior was not acceptable.

A 'timing' problem was thought of. A timer is an internal clock that 'ticks' very often, in the Honeywell, one million times per second. It synchronizes all the operations of the computer. If a packet arrived at just the wrong time between ticks, the computer would go into a state of 'shock,' and would crash. But, the group found the trouble and the machine ran for a day with no crash. But, there was no time to make a multi-day test, because the IMP was to be shipped the next day.

The IMP was crated and delivered to the airline, with precautions written all over it for delicate handling. Barker traveled at the same time on a separate passenger airplane. When it arrived at UCLA everyone was amazed by its size, that of a refrigerator, weighing nearly a half-ton, encased in drop-tested, battle-ship-gray steel.

UCLA needed electronic hardware to connect their host computer to IMP No. 1. One graduate student, Mike Wingfield, volunteered to build the hardware. He had some experience, already having built a graphics interface for another computer. He miraculously completed the task a week before the IMP arrived. Soon the IMP and the host computer were 'talking' to each other flawlessly. On October 1, 1969 the second IMP arrived at SRI, exactly on schedule. This delivery marked the first possibility of a real ARPANET test. The two hosts, as with all later ones, were connected together by a leased 50 kilobit telephone line. The story of the first message was told at the beginning of this article.

A very important ingredient to the successful use of the ARPANET was missing. There was no host-to-host protocol, that is to say, there was no set of operating instructions that would let the disparate computers 'talk' to each other. The minimum duties the host-to-host system should be able to perform were: (1) remote logins, i.e., allow the user at host 'A' to connect to the computer at host 'B' and work with it as though the 'B' computer was at site 'A'; (2) transfer of files (documents, spreadsheets, etc.) from one site to another; and (3) ability to process, i.e., create or edit, files at another host. It would be over a year before a minimum host-to-host protocol was created.

The inter-host software was called the Network Control Protocol. It was hammered out by the Network Working Group, a joining of the staffs of the host sites and BBN. The process was difficult and there were many opinions and suggestions along the way. With the Network Control Protocol available, the ARPANET architects could now boast that the entire enterprise was a success, packet switching unequivocally provided efficient use of communication lines. It was clearly an economical and reliable alternate to circuit switching, the basis for the Bell Telephone system.

**Tomlinson is best known today for having selected '@' as the division between a recipients name and his/her location.**

Little by little BBN was developing a center for the maintenance and control of the network. This evolved into a large room equipped with computers and teletypewriters which were programmed to sample the IMPs, minute by minute, via the interconnecting lines to determine if they were operating properly. Hourly summaries of the reports were compiled. If no status report was received from an IMP for three minutes, the teletypewriters would show which IMP had crashed and why. As the network grew, a large map (a metal sheet with magnets) was produced on one wall of the control room that showed the wiring of all the IMPs, host computers and links. Before long, a network maintenance team was set up at BBN to respond quickly and authoritatively to any failure.

Even by the end of 1971, the ARPANET was little used because the hosts that were plugged into it were not all equipped with the necessary software. Janet Abbate [*Inventing the Internet*, Cambridge, MA 1999] summarizes, "*The obstacle was the enormous effort it took to connect a host to an IMP. Operators of a host had to build a special-purpose hardware interface between their computer and its IMP, which could take from 6 to 12 months. They also needed to implement the host and network protocols, a job that required up to 12 man-months of programming, and they had to make these protocols work with the rest of the computer's operating system. Finally, they had to adjust the applications developed for local use so they could be accessed over the network.*" Thus it appeared that people needed to be inspired to use the ARPANET.

## E-mail

It was in 1972 that the first real electronic mail was delivered. Before then, a communication could be written into a file folder and other people on the same computer could ask for transfer of the folder. After what was really a 'hack' between two computers at BBN, Ray Tomlinson at BBN wrote a mail program that had two parts, the one to send was called SNDMSG and the other to receive was called READMAIL. Tomlinson is best known today for having selected '@' as the division between the recipient's name and his/her location. What a success – by 1973, three quarters of all traffic on the ARPANET was e-mail. It was remarked, "You know, everyone really uses this thing for electronic mail." Larry

Roberts further streamlined e-mail by writing a program for listing the messages and a simple means for accessing and deleting them. A further valuable contribution was 'Reply,' added by John Vittal [at BBN], which meant that the whole address did not need to be retyped when answering a message.

### The Internet

When ARPANET was planned, no one envisaged that other networks would need to be addressed. One IMP could only communicate directly with one or two other IMPs. And the packets were created by the IMPs and reassembled by them. The hosts delivered complete typewritten messages (and graphics) to them and received only reassembled messages. Thus, the only address needed for a packet was that of the destination IMP, because the actual computer or person for whom it was intended could be named in the message itself.

By the early 1970s other packet-switching networks were being created. If a message from an outside network was to be sent to the ARPANET, the packets would already have been created by the former. Also, a meandering packet would be addressed to any one of a number of networks, so it must bear both the name of the destination network and the destination host on that network. Customarily, also, the packet would name its source network and host so that an acknowledgment could be made.

The interconnection of separate networks lead to the next major stage in the unfolding of the Internet. Robert Kahn and Vinton Cerf generated a new protocol (set of rules), for marking internet packets so they would be delivered to the proper destination. This protocol was called TCP/IP and was published in 1974 ["A protocol for Packet Network Intercommunication," IEEE Transactions on Communications, May 1974, 637-648]. The IP part of the Kahn/Cerf protocol was rules for the header of each packet that, like the front of a mail envelope, contained the destination and return addresses. The TCP portion comprised the rules for breaking up the message into packets, reassembling them at the destination in proper sequence, making a checksum and resending any packets that did not arrive or were deficient.

### The Ethernet

Today's chosen electronic equipment for breaking up a message into packets and reassembling them and for supplying the header was invented by Bob Metcalfe and David Boggs. The Ethernet differed from the ARPANET in several important respects. First, the Ethernet card splits a message, typewritten or graphics, into packets. Second, it is a broadcast system, that is to say, the packets are introduced onto the Ethernet cable by any node and are sent by it to all other nodes simultaneously. Only the node to which they are addressed accepts them, the others reject them. Finally, at its destination, the Ethernet card reassembles them. In the ARPANET system, the packets only traveled from one IMP to the next where they were stored for a microsecond and examined to see if they are for that IMP. If not, they were forwarded to the next IMP and the next, until they finally reached their destination. The broadcast feature of the Ethernet reduces the cost of each node and lowers the transit time on the cable. The packets can be of variable length, and there are procedures for handling 'collisions' between two packets and two packets are prevented from being introduced onto the cable at the same instant.

On January 1, 1983, ARPANET made its official transformation to TCP/IP. That is the official date of the formation of the INTERNET, the word that signifies the collection of all networks.

In 1983, there were 562 nodes on the ARPANET. The government decided that ARPANET was so large that security was no longer possible. It was split into two parts: MILNET for the government laboratories and ARPANET for the others. In 1985, the Congress created five supercomputer centers in the U. S. Agitation grew for the government to create high capacity lines, called the 'backbone,' to join the super-computer centers together. The backbone was established in 1985, and in 1988, it was upgraded to a speed thirty times that of the ARPANET. A few years afterward, it was upgraded to 1000 times ARPANET's speed. Another significant network was added in 1991, financed by the U. S. High Performance Computing Act, sponsored by Vice President Al Gore. It was called the National Research and Education Network, NREN.

With privately-supported networks popping up everywhere, the government decided that it could eliminate the expense of maintaining ARPANET, which would save about $14 million per year. The nodes on it were migrated to nearby networks and decommissioning was accomplished by late 1989, just twenty years after its first 'ello.'

All indications are that early in the 21st Century, an equal number of households on the globe will be connected to the world wide web as have television today. Only a few years ago personal computers, PCs and Macintosh computers, were in homes and offices only for utilitarian reasons. Anyone who had to write documents and those maintaining financial data found word-processing and spread-sheet programs invaluable. Besides the WP and spread-sheet programs the 'office' programs also have a calendar for appointments, a card file, an address book and a relational data-base program. Photos and printed matter can be converted to digits and put in memory by scanners. Wide choices of fonts and editing tools make possible the production of camera-ready pages, with illustrations, for sending to printers of books.

Computers with high speed chips, large random access memories and disk storage memories of humongous size make these applications a pleasure to use. The change from serious to personal use began a few years back when fast moving games for rendition on a computer could be bought on compact disks, CDs. But, mostly children took to those. Suddenly, since about 1994, there are web pages, search engines, e-mail, news and movies, all available, through modems, via the local telephone and high-speed cable systems. The Internet has succeeded beyond expectations, partly because it is 'fun.'

### The world wide web www

The story of the world wide web is fascinating. It was only through the persistent struggles of one man that the entire world got its act together to make 'www' today's reality. Tim Berners-Lee, fashioned the protocols for the World Wide Web gradually between the years 1980 and 1991. He was stationed at CERN, the International European Particle Physics Laboratory in Geneva, Switzerland. [*Weaving the Web: The Original Design and Ultimate Destiny of the WORLD WIDE WEB, by Its Inventor*, (New York, 1999)]. The challenge Berners-Lee confronted was, in his words, "The computers simply could not communicate with each other." Fortunately, the basic ingredients for the world wide web had been invented, available for him to combine, season and bake into the next marvel. These basics were Baran's packet switching, BBN's and the cooperating institutions' ARPANET, Bob Kahn's and Vint Cerf's TCP/IP, Douglas Englebart's 'mouse' (1960), and Ted Nelson's 'hypertext' (1965). The latter two, play a vital part in the WWW.

The short time in which such Internet services have been attainable, starting about January 1995, is inconceivable. The URLs that result from searches or lists, make available: titles and often the pages of books in a library, the complete contents of a half-dozen encyclopedias, the home pages of thousands of companies, institutions and individuals, medical advice and hundreds of other categories.

Many people have claimed to have invented the ARPANET. Roberts, put together the "request for proposal," after consulting others, and he could claim invention of the idea. BBN, in the language of the patent office, "reduced to practice" the concept of a packet-switched wide-area network. Stephen Segaller [*Nerds, 2.0.1*, (1998),182] wrote that, "What BBN did invent was *doing* packet switching, rather then proposing and hypothesizing packet switching" (emphasis in original). Frank Heart states: "We were unable to use any of the work of Kleinrock or Baran in the design of the ARPANET. We had to develop the operating features of the ARPANET ourselves." A detailed account of BBN's effort is given by Hafner and Lyon, *Where Wizards Stay Up Late*, (1996) Simon & Schuster, New York.

S/V

---

The author may be reached at: beranekleo@ieee.org.