Len Kleinrock: The First Two Packets on the Internet

Charles Severance, *University of Michigan*

Len Kleinrock describes the Internet's humble beginnings.

Although it might seem like the Internet has been around since the beginning of time, it did indeed start somewhere. The first two packets ever sent on the Arpanet originated in a lab at the University of California, Los Angeles, and ended up at the Stanford Research Institute on 29 October 1969.

I visited with Len Kleinrock at UCLA to learn the story of those first two packets; visit www.computer.org/computingconversations to view our discussion.

OFF TO THE RACES

ARPA funded the Arpanet, the first version of the Internet, because it wanted to make the best use of its research investments, which were scattered in various university campuses across the country:

> The University of Utah had a terrific graphics operating system, SRI had database expertise, we had simulation technology, and the University of Illinois had high-performance computing. Every time ARPA brought on a new researcher, it offered to buy a computer for that person, but researchers invariably said, "I want the same capabilities that all those other researchers have: graphics, database, and all the rest." ARPA said it couldn't afford that and offered access to those resources through a network rather than by replicating them; the need for the network was for resource sharing, not to protect the US against a nuclear holocaust.

Building a multi-institutional, wide-area network to connect diverse computing hardware would be a significant effort that needed careful coordination:

> ARPA brought in Larry Roberts—an office mate of mine at MIT—to manage this project. He came to me because he knew my work; he watched me develop my simulation using his compiler on the TX-2 computer. He said, "Len, we need to know if this thing is going to work." He knew that I had the theory, so I could show him that it would work. Later, he said that he never would have decided to spend millions of dollars of the US government's money if he wasn't sure that it would work.

Len and Larry spent about a year bringing experts together to develop a specification for the new network's software and hardware:

> During the design phase, some great people joined us. Herb Baskin said, "If this network can't deliver short messages within a half-second, I can't use it for time sharing," so that became part of the specification. We got it down to 200 milliseconds. Wes Clark said the switch needed to be a completely separate computer for communications. Artificial traffic generators, measurement hooks, a way to evaluate the measurements—as a network researcher, I put all of that into the requirements.

The design was developed throughout 1967, and in 1968, the team sent out a formal Request for Proposal. Bolt, Beranek, and Newman (BBN) won the contract to produce the Arpanet's first switch, and UCLA functioned as the first node and the Network Measurement Center, responsible for testing and verifying the new networking hardware and software:

> All those specs went to BBN, which built the darned thing and delivered the first switch to UCLA eight months after it got the contract, on time and on budget. It came here, we plugged it in, and bits began to move back and forth

Published by the IEEE Computer Society

> between our time-sharing machine and that switch on the day after Labor Day, 2 September 1969.

LO AND BEHOLD

The second switch went to the SRI in Palo Alto in October 1969, and once it was connected to the SRI host computer, the team could join the two switches and test them. The initial 50-Kbps connection was made up of multiple 4.8-Kbps leased lines ganged together:

> One night in late October, Charley Kline and I said, "Let's communicate between these two machines," so we got hold of Bill Duvall at SRI and said, "Let's log in from a terminal connected to the UCLA computer and then from the UCLA computer, log into the SRI computer over the network." The host computers were both time-sharing systems that expected terminals to connect locally and use the machine's services. The plan was to sit at a UCLA terminal and log on to the UCLA machine and then through this wonderful network, log on to the SRI computer as if we were a local user at SRI.

To help debug the new software, the team had an analog audio connection using one of the 4.8-Kbps lines between UCLA and SRI. They watched both systems and the communication channels very carefully as they started the test:

> Charley typed the L, and we asked if Bill saw it—he said yes. Charley then typed the O, and Bill saw it on his end as well. We were trying to do "LOG" for "LOGIN," but when Charley typed the G, the SRI computer crashed. So the first message ever on the network was "LO" as in "Lo and Behold."

For all the years since 1969, Kleinrock has kept interface message processor (IMP) serial number 1 at UCLA. There has been discussion of putting it in the Smithsonian, but he wanted to ensure that the IMP was always on display for the public to see. It currently sits in Boelter Hall in the exact room where those first two packets were sent from UCLA to SRI. The room has been restored to look exactly as it did in 1969. Even the architects at UCLA are finding ways to commemorate the historic events that happened in Boelter Hall that day:

> If you come in the back entrance to the building, you walk on a mosaic of tiles. It has a strange pattern that turns out to be ASCII code for "Lo and behold!" The tiles memorializing our message were quietly installed there by the UCLA architect, Erik Hagen, about a year and a half ago now, only to be recently discovered by Viet Nguyen, a first-year computer science undergraduate student.

THE NEXT GENERATION

Although UCLA was initially responsible only for testing and measuring the Arpanet, over the next decade, it assumed leadership of the Arpanet effort as graduate students including Vint Cerf, Steve Crocker, Jon Postel, and Charley Kline became the project's leaders, managers, and developers. Initially, BBN saw the Arpanet and IMPs as "black-box" products, but ARPA ultimately forced BBN to open the system's source code:

> We could break it at will, and every time we did, we would call BBN and tell those guys to fix it because they wouldn't give us the code. They kept it proprietary until ARPA said that it had paid for the code and forced BBN to open it up. Once we had the code, we would discover a fault and would show them how to fix it, but it still took them six months to fix it.

Over time, the project moved to a more open source approach overall and drew its technical directions and leadership from the growing Arpanet community:

> One of the things I was very much interested in was distributed control. I was a student of Claude Shannon, and his best work came when a lot of things interacted—long code words, for example. That's when these emergent properties would arise. For this reason, I wanted to design large networks, and to design a large network, you cannot have a single point of control—you need to distribute control. When you distribute control, you delegate authority to all peers.

ARPA believed in the researchers involved in the project and gave them a lot of autonomy:

> When ARPA started funding principal investigators, it had the same philosophy of delegating authority—"You're a smart guy, here's some money, go do the thing you do best, and we won't sit on top of you." When I received that kind of money, I did the same with my graduate students. They're brilliant kids, so I would tell them, "We need a host-to-host protocol. You build it—I'm not going to sit on top of you." That isn't a "product" mentality—that's a research and development and creative mentality. And it works so well.

It certainly did work well, and after nearly 45 years, the notion of distributed control, emergent properties from large networks, and delegating authority to our peers makes it possible for us to almost forget that perhaps the most complex collective engineering endeavor humanity has ever attempted works so well we barely notice it anymore. "Lo and behold," indeed.

Charles Severance, *Computing Conversations column editor and* Computer*'s multimedia editor, is a clinical associate professor and teaches in the School of Information at the University of Michigan. Follow him on Twitter @drchuck or contact him at csev@umich.edu.*