



DOI:10.1145/3591217

Leah Hoffmann

Q&A

Still Networking

Bob Metcalfe, who co-invented Ethernet networking technology, continues to take on new challenges as a computational engineer.


ACM A.M. TURING Award recipient Bob Metcalfe—engineer, entrepreneur, and Professor Emeritus at the University of Texas at Austin—is embarking on his sixth career, as a Computational Engineer at the Massachusetts Institute of Technology (MIT) Computer Science and Artificial Intelligence Laboratory (CSAIL). He is always willing to tell the story of his first career, as a researcher at the Xerox Palo Alto Research Center (PARC) where, in 1973, Metcalfe and then-graduate student David Boggs invented Ethernet, a standard for connecting computers over short distances. In the ensuing years, thanks in no small part to Metcalfe's entrepreneurship and advocacy, Ethernet has become the industry standard for local area networks.

Leah Hoffmann spoke to Metcalfe about the development of Ethernet and what it has meant for the future of connectivity.

You published your first paper about Ethernet in *Communications* in July 1976 (<https://bit.ly/403Sxmm>). What do you remember now about the publication process?

Well, I wrote the paper in 1974, but Xerox wouldn't let me publish it, because we had to patent everything first. Then there was a legal review. So the paper benefited because it got rewritten 20 times, and I loved the paper because they let me write it.

You borrowed the name “ether” from 19th-century physics. Then, it was thought to be a passive, omnipresent medium for the propagation of light



and gravity. In your conception, Ethernet would be an omnipresent medium for the propagation of packets of electromagnetic signals. The idea is perfectly suited to a post-Wi-Fi world, but how did you come to it at a time when those

“Because we had started with radio and then begrudgingly went to cable, that led into the name ‘Ethernet.’”

packets were propagated not through the air, but through coaxial cable?

Well, we didn't start with cable; we wanted no wires. And then I went to Hawaii to see the ALOHAnet. I got two things from Hawaii. First was this randomized retransmission technique for taking turns on the channel.

This is the idea—which first reached you via a 1970 National Computer Conference paper by Norm Abramson about ALOHAnet—that if two stations interfere with one another by trying to send at the same time, they should try again at some randomized interval in the future.

That's right. The other thing I saw in Hawaii was an ALOHAnet modem for driving the radio. That modem was bigger than [CONTINUED ON P. 118]



[CONTINUED FROM P. 120] the Alto, and the Ethernet had to fit on a card. So, that was my second takeaway from the trip: radio is a non-starter. But because we had started with radio and then begrudgingly went to cable, that led into the name 'Ethernet.' The original name of Ethernet was actually the Alto-ALOA Network, because it connected Altos and used ALOHA for taking turns. It was only in the so-called invention memo (<http://bit.ly/406xDDi>) that we changed the name to 'Ether.'

I realize you have been telling this story for a long time, but can you summarize what stands out for you now about the technical decisions you made while building Ethernet?

There were three technological decisions that we all thought were important at the time. Now, 50 years later, there are three other things that I see as really important.

Let's start with the original list.

Here were the three things that we built Ethernet around. One of them was coaxial cable. We chose it because there was a device in the cable called the Jerrold tap, thanks to my friend David Liddle, who came from the cable industry. Through the Jerrold tap, you could tap into the coaxial cable without cutting it. That meant we could keep the network

up while we were installing a new PC.

Number two, we chose an encoding technique called Manchester encoding. That's the transceiver—the thing that drives the cable by taking it up to several volts, then letting go. Manchester is the simplest possible encoding. In each bit's cell, the transceiver is on half the time. If it's a one, it's on and then off. If it's a zero, it's off and then on.

It had other features, too. Using Manchester meant that, if you were sending a packet, the network was off half the time—so if you were not transmitting this very instant and someone else was, you could see it. You could do collision detection. The other fea-

"Ethernet brought packets to the desk. Before that, only characters got to the desk, which limited what you could do."

ture of Manchester is that it has an implicit clock in the middle of every bit's cell. Each bit either goes up to down or down to up, and that's the ticking of the clock. That meant that Ethernet did not need a central clock, because every packet had the clock in it during these transitions.

Third, we needed a way for the stations to take turns. That brings us back to randomized transmission. So, randomized transmission, the Jerrold Tap, the Manchester Encoding—that was Ethernet.

And now the list is different.

But it still has three items on it. Three is my favorite number. So, first of all, Ethernet brought packets to the desk. Before that, only characters got to the desk, which limited what you could do. The TELNET Terminal Protocol ran in big time-sharing machines instead of on your little PC. Ethernet changed that.

Two, the packets came really fast. Ethernet began the tradition of abundance. It was 10,000 times faster than what it replaced—not ten, 10,000.

Third, Ethernet joined this novel thing of standardizing. The Ethernet standard went right into the IEEE 802, which did a fabulous job of standardizing it. In hindsight, those three things mattered more than the three technology decisions.



What do you think about when you think about Ethernet's future?

Well, it raises the question, what does the word Ethernet mean? It's been 50 years, so it's had a chance to mean a bunch of different things. Today's Ethernet is not the one that David Boggs—may he rest in peace—and I built in 1973. It's completely different. But some invariants have emerged. One is that Ethernet is an open standard. Another feature is it's faster than anybody needs. It's interoperable. There was a big debate back then about certifying things to be standard. Fine, meet the standard, but you better be interoperable or you're not going to last.

But going back to your question, Ethernet was the transition. The Internet began by connecting time-sharing machines, principally PDP 10s. Then along came the personal computer. The first thing I wrote at Xerox was a TELNET for the Alto—so you could sit at your Alto and log into PDP 10s, as if that were interesting. The personal computer and the Ethernet both happened at exactly the right time. Ethernet was the technology that enabled the transition of the Internet from dumb terminals to personal computers. Now, the word Ethernet is a brand. It's a promise to its users about things like interoperability and speed and standardization.

I'm going to re-ask a question that Robert Colburn asked you nearly 20 years ago (<http://bit.ly/3n7wSet>). What exciting or rewarding or unexpected ways have people found of using the Ethernet, in your view?

Twenty years ago? Time flies. I'll start with automotive Ethernet, which is gratifying because it's a new space. Some of the people who build cars are planning to put Ethernet in them. And I've detected that there's already a fight between different ways of doing it. It's very familiar scene. People recognize the need for standardization, but there's no agreement on

what the standard should be. Which has always been the normal situation of Ethernet.

IoT is also a huge space. And there's a group called the Ethernet Alliance that seems to be doing the right thing. We used to call it co-opetition—you cooperate on the standard and then you compete like hell. That's been typical of the Ethernet. Ethernet was protected by patents, which were put into the public domain. And then my company licensed it, and suddenly we had 100 competitors.

Your next career move takes you back to MIT.

After I retired from the University of Texas, they made me emeritus—which I appreciate, although nobody knows what being 'emeritus' means. Then, I decided my new career would be computational engineer. So I went out and told everybody, and my alma mater, MIT, said, "Why don't you come and do your computational engineering here?" I leapt at the chance. Now I'm a research affiliate of the Computer Science and Artificial Intelligence Laboratory. My first engagement is simulating geothermal wells.

Leah Hoffmann is a technology writer based in Piermont, NY, USA.

© 2023 ACM 0001-0782/23/6 \$15.00

"Today's Ethernet is not the one that David Boggs—may he rest in peace—and I built in 1973. It's completely different, but some invariants have emerged."