

<http://www.telecomkh.com/en/business-communications/news/ethernet-innovation-summit/netevents-global-press-analyst-summit/5396>

Meet the inventors of Ethernet

Date: Thu, 05/30/2013

Celebrating 40 years of Ethernet Innovation at the Computer History Museum, Mountain View, California



From left to right, Dave Boggs, Norm Abramson, Bill Howe and Radia Perlman

Image credited to NetEvents

A reunion of the brains behind the biggest revolution in human communications since the invention of movable type – and a rare opportunity to see them reflect on those times. What was the world like before personal computers? How did Ethernet emerge from Arpanet and ALOHA? What was the need? Had they any idea where this invention would take them, and the world?

Panellists: Norm Abramson, ALOHA; Dave Boggs, Ethernet Co-Inventor; Bill Howe, DEC ; Radia Perlman, Spanning Tree Protocol Inventor

Chair: Paul Saffo, Futurist

Our first panellists up, Norm Abramson and Dave Boggs and Bill Howe and Radia -- Radia, sorry. I'm going to get this right. I have another friend called Radia -- Radia Perlman. Radia. Radia, I swear I will get this right eventually. Come on up and grab a seat. So this panel, as we start here, this is "Meet the Inventors of the Ethernet." Can you put me down on this side and just shift everybody

down one? That's good. And the way we're going to do this, so we have -- Radia, you're right next to me. Fabulous. We have 30 minutes. We have 40 minutes? I'm going to shift this so I can see it from over there. Our time is short. You're going to learn more about each of these individuals in detail. We agreed in advance that the way I would introduce people is with haiku, rather than go down the lengthy careers and many distinguished accomplishments of them.

Radia, 10 base T. Norm Abramson.

Unidentified Participant

Wrong.

Radia Perlman

Spanning Tree.

Paul Saffo

Spanning Tree, sorry.

Unidentified Participant

It has a T in it.

Paul Saffo

Sorry, Spanning Tree. Norm Abramson, ALOHA, Dave Boggs, co-inventor. Can we leave it at that?

Dave Boggs

Yeah.

Paul Saffo

Okay. And Bill Hawe, [Disp] Internet. So let's get started, here. I want to start with you, Dave, co-inventor. What exactly does that mean? Tell us what it was like to be co-inventor.

Dave Boggs

Gosh, it was Bob's idea. I just helped him reduce it to practice, as they say. I was a grad student at Stanford, and every Friday afternoon, the master's candidates were required to attend a lecture from outside people, and they took attendance. Friday afternoon, mandatory lecture at 3:00, and they took attendance. So, one time, Alan Kay gave a talk. He mostly talked about Smalltalk, and that pretty much went right over my head. But at the end, he showed a bunch of pictures of six-foot-tall equipment racks full of blowers and cards, and he said they were building a PDP-10 clone. So I was exploring doing PhD research in high-performance CPU design, and that caught my attention. So I pestered Alan for a job at the Learning Research Group, and he looked at my resume, which was I'd worked at NBC as a broadcast engineer and built a lot of hardware, so he passed my resume on to another group that was doing a clone of Doug Engelbart's video terminal. What's the name of Doug's system?

Unidentified Participant

NLS Augment.

Dave Boggs

NLS, yes, right, at SRI. So because I had background in television, I ended up in that group, fixing the televisions. And the system was based on 15 Data General Nova copmputers that were all lashed together with a Data General proprietary network, time division multiple access system, called the Multiprocessor Communications Adaptor. So with that many computers, one of them was always broken, and I ended up being the keeper of the [test-and] machine, where you'd swap parts out, and occasionally somebody from Data General would come out and use that to fix them. So, one day, Bob had designed an Arpanet interface for the Xerox clone of a PDP-10, which was front-ended by Data General Nova, so his Arpanet interface plugged into the Data General Nova, and he built it and it worked. Then they wanted to build another copy of the PDP-10 clone, and so they needed another Arpanet interface. Well, Bob's drawings weren't really quite up to date, and he wasn't really sure he could build the second one, so he them wire it up with whatever he had, and then he needed a machine to debug it in, so he came to me, because I had the [test-and] machine. So we spent a

couple of weeks figuring out all the changes that he'd put into the first Arpanet interface that he never wrote down in the schematics. And then, little did I know at the time, he was writing all of those memos, inventing Ethernet, and so he started explaining it to me. And the pressure was growing on him to stop writing memos and start building hardware, so he'd been working with me, he got along with me, I guess, so he turned to me and said, "Would you like to help me build this idea?" And he handed me a stack of memos, and that's how I heard about Ethernet.

Paul Saffo

Fabulous. Radia, you look like you're just about to ask a question.

Radia Perlman

No, no.

Paul Saffo

Well, I have a question for you. So I know that there are some who call Radia the mother of the Internet, which here in California, we all share in her wincing at because it is so politically incorrect to put. And I would actually invite the audience, we need a better phrase for this, because father of the Internet sounds okay. Mother of the Internet doesn't quite work, so we need a gender-neutral name. This would be a good innovation for the industry to come up with. But leaving that aside, just how did the Internet evolve?

Radia Perlman

Okay, well, actually, I'll talk sort of about the Spanning Tree piece of Ethernet. But, first, I'll say that, gee, in keeping with this event, this protocol for asking questions, raising your hand, shouldn't be like CSMA/CD, that you just?

Paul Saffo

Awesome. Hand raising is over. Just barge in at any time, and we'll do the error correction up here.

Radia Perlman

Right, and the other comment about CSMA/CD is that I seem to have it ingrained in me. That's the way I behave. I cannot interrupt somebody, and if somebody interrupts me, I just stop, but I'm always amused at meetings when there's other people that don't seem to understand CSMA/CD, and whenever they feel like talking, they just talk, even if someone else is talking. And they'll keep talking.

Paul Saffo

That's actually most software engineers I know.

Radia Perlman

Right, right. So back to Spanning Tree, I was the layer-3 designer for DECnet, which meant that it's how you piece together a bunch of links and have the network self-forming and figure out how to move the data around with intelligent switches. That's another way to hook up a bunch of computers, and all the links were like point to point. It was only two computers. Then, along came the Ethernet, and it was like, "Oh, gee, a new kind of link that I have to redesign the routing protocol to be efficient with this new kind of link." So I applaud the inventors of Ethernet. It's really cool technology, but they should have called it Etherlink and not Ethernet, because it was all bright and shiny, and everyone got all kind of confused, and they thought it was instead of layer 3. So they were building applications directly on Ethernet and leaving out layer 3. Ethernet was intended to be just a single link, with a limited number of computers and distance that you could do. So I tried to argue with them and say, "Hey, hey, you still need layer 3."

Unidentified Participant

How'd that work out?

Radia Perlman

Right, right. They said, "Oh, go away, Radia. You're just upset because no one needs your layer anymore." And I said, "But you may want to talk from one Ethernet to another." And they said, "Our customers would never want to do that." Years later, my manager says, "Gee, we need to figure out a

way, a magic box, that will sit between two Ethernets and let someone over here talk to somebody over there." But, of course, the constraint was we couldn't change the Ethernet packet in any way. We couldn't increase its length. There was not a single spare bit, and we couldn't change the way the end nodes worked. There were end nodes that thought they were speaking on a single link. They had pretty much agreed that, well, you just kind of listen promiscuously and forward the Ethernet packets around, but there needed to be a Spanning Tree algorithm, which spanning means reaches everybody, and tree means no loops, because that wouldn't work, if there was multiple ways to do that. My manager suggested this problem on a Friday, and he was going to be gone all the next week. I realised that night how to do it, and it was trivial, and it was incredible. I could prove it worked and all that. I spent Monday and Tuesday writing the spec in enough detail that people got it working without even asking me a single question in just a couple of months. And the remainder of the week, because I couldn't concentrate on anything else -- I had to show off to my manager, and he was away, and this was before cell phones or e-mail or electricity or whatever. So I spent the remainder of the week working on the poem that goes along with the algorithm, which was the abstract of the paper in which it was published.

Bill Hawe

There was some controversy at that time about whether this should or shouldn't be included in the bridge, because the hardware development team's default position was no extra stuff, period, and it's not a bad position for a hardware team to say, "I want it to be as simple as possible, but no simpler." And so they said, "We don't have it in repeaters, so we don't need it in bridges."

Radia Perlman

And that's another actually really interesting story, which was that I found out afterwards that the very first bridge that Digital sold was to the world's most sophisticated customer. I found out afterwards that they were saying, "Oh, we're doing all of this complicated stuff," and the Digital salesperson was saying, "Eh, it's just going to work." And they were saying, "But we're doing this complicated stuff." So the world's most sophisticated networking customer had the world's simplest topology, which was two Ethernets and one bridge, and they plugged it together and it didn't work, and they were really upset. And when field service went to figure out what had happened, they realised this sophisticated customer had plugged both ends of the bridge into the same Ethernet cable. And so everything was working perfectly. The Spanning Tree algorithm said, "Hey, I don't need to forward. If I ever do, I will."

Bill Hawe

Yeah, a related disaster, and there were certainly many disasters along the way that got fixed, before the first bridge was sold, was we said, "Well, let's resolve this question about whether there should be a routing algorithm in the bridges." And so down in our lab in Hudson, Massachusetts, where corporate research had moved, we built one with a PDP-11 and a couple of Unibus adaptors, and it was way bigger than that podium. And hooked it up, again, on a Friday, went away, and then sometime over the weekend, somebody in there created a loop, plugged it in, and it brought down the entire Hudson facility. That's now an Intel semiconductor facility. At the time, they were running a distributed simulations for the development of the MicroVAX computer, using every computer in the building. And Bob Supnik's head exploded at that point, and it landed directly in front of me, and it was yelling at me. And so that got pulled off, and he said, you ought to just pull corporate research off of this network, as well, beside this experimental thing. And that resolved for the hardware team whether or not there was going to be a routing algorithm in the bridges, definitively.

Paul Saffo

Radia, I just want to go back to one thing that caught my attention, what you said. You said, "So I went home that evening and the solution occurred to me," and then you kind of moved on. It's like, "Well, so, then magic happened." What happened? What was the insight?

Radia Perlman

Oh, it's after the fact you look at it and it's just sort of obvious how to do this incredibly simple thing. Should I quote the poem?

Paul Saffo

Please.

Radia Perlman

Okay. So the poem is called Algorhyme, because every algorithm should have an algorhyme, and it's, "I think that I shall never see a graph more lovely than a tree. A tree whose crucial property is loop-free connectivity. A tree which must be sure to span so packets can reach every LAN. First, the root must be selected. By ID, it is elected. Least-cost paths from root are traced. In the tree, these paths are placed. A mesh is made by folks like me, then bridges find a spanning tree."

Paul Saffo

And, by the way, this will be the sobriety test after cocktail hour tonight. Anybody who cannot recite that completely will not be allowed to drive home. They'll set up cots here in the Computer History Museum. So, Norm, you're off in Hawaii, tooling along. You're a professor doing your research. At what point did you begin to suspect there was something weird going on in Silicon Valley?

Norm Abramson

Well, I got to Hawaii by spending 10 years in Silicon Valley, so I knew long before I got to Hawaii that there was something weird in Silicon Valley. I don't know when that struck Bob, but Bob was, as he usually is, very gracious in crediting the work in ALOHA to what he brought to fruition in so many other places.

Paul Saffo

Well, he was looking for that chapter in his thesis. He had contacted you, of course.

Norm Abramson

No. He only contacted me when he finally got to Hawaii, and in fact, the whole process by which Bob got to Hawaii is one of these things that's sort of shrouded in the mysteries of networking. I don't know how it occurred. Bob says I invited him, from time to time.

Paul Saffo

I thought it had something to do with the statute of limitations and jurisdiction.

Norm Abramson

It may have to do with the passage of time, but I don't think there's any statute involved. The first I knew is that Bob showed up as perhaps the only predoc postdoc I've ever known. That is, Bob had not got his degree, I believe, when he got to Hawaii. He was still fighting with the people at Harvard, and the work that he did in Hawaii was helpful to us. Bob was involved with a lot of our graduate students, and Bob and I talked from time to time, but it was networking from the point of view of the ALOHAnet. And perhaps, one of the things that I should bring up is the question that was mentioned just peripherally in Bob's statement here. He's talked about the path from ALOHA to Ethernet. Of course, if you take a look around at all of the computers that you see at the various desks here, everybody's connected not by Ethernet but by wi-fi, and so there is a road from Ethernet to ALOHA that is interesting to me, as well. And it seems to be the way in which things are going, and with a great deal of pleasure I notice that the movement to the wireless end of what's going on in Ethernet seems a lot more robust and also a lot more lasting, I think, than the wired version of ALOHA which Bob brought out. I certainly see Bob as having invented both.

Paul Saffo

So instead of ALOHA being the ur-Ethernet, Ethernet is the ur-ALOHA, and it's only appropriate that ALOHA means both hello and goodbye, so 40 years ago you said hello to the Ethernet. Now it's like, "Well, so long. It's ALOHA's future."

Norm Abramson

And we really haven't said goodbye to the Ethernet, but we have said hello to ALOHA again, after a hiatus of perhaps 20 years. That hiatus, by the way, was not a technology hiatus. It was a regulatory hiatus. Until the frequencies were available for the connections that we all use today, there was no sense in going from Ethernet to wi-fi or ALOHA.

Paul Saffo

So, Bill, you of course have been in the thick of all this. One thing I should share is that we all had this really wonderful e-mail exchange before this panel, and I have got like 10 pages of really cool stories. Bill told some of the best, and I don't even know where to begin, but Bill, what was the part that you thought was most likely to fail in all of this?

Bill Hawe

Well, there was a whole list of things, but because it's much easier on the ladder of credibility to sort of fall off at any point than to actually get to the top and commercialization. For example, Jerry Saltzer at MIT had written a paper, "Why a Ring," which was all of the arguments for token rings. This is back in the sort of early-1980 timeframe, when we were building the 10-megabit Ethernet. And I put it on my cubicle, that we were in the Mill at DEC at the time, with a dart, the kind you use in a bar. I just put it there and made sure you went down that list. He said a few things like, "You'll never get the random number generators to work." And so we thought, "Oh, we can do that," and so we had a design, we got it going and everything, and we built, much to the annoyance of Intel at the time, because with the Digital, Intel, Xerox, I think there was an implicit assumption that we were just going to use the Intel chip. So in order to fulfil that assumption, we designed user own chip and then licensed it to AMD and MOS Tech, figuring between the three of them, somebody would get it right, and that would drive the costs down and the bugs out and create a dynamic. The commercial terms on it were tremendous in terms of what we would have to pay. And so we designed that. We built three emulators of the chip and ran it, and we were trying to make sure everything really worked for commercialization. And over the weekend, all the systems drifted into sync, and so the collisions were -- [backoffs] were permanent. There were zero packets going through the system. It would lock up. The closer the crystals were in frequency, the longer it would take to drift into sync and the longer they would be in sync, once the random number generators were synchronized. So this created a fire drill. Bernie Lacroute, who was sort of the vision manager before he had left to go to Sun -- we called him the Crusher. I'll tell you offline why.

Paul Saffo

No, just TTY interrupt, I think you need to explain Crusher.

Bill Hawe

Well, because a couple of us had gone and wanted a little bit more time, because I had found some issues in the backoff algorithm in these simulations we were running, this could be fixed and so forth, and so we just needed a little more time. One person made the mistake of saying, "Bernie, you can't get blood from a rock." And he says, "No, but I can crush a rock," and he just crumpled up the paper into it, so he became called the Crusher. So the things drifted into sync, and so we were all locked in a virtual room by Bernie and management to fix that problem, and the reason we wanted it fixed, and this relates to I think one of the very important commercial dynamics that was established with the 10 megabit, which was establishing a dynamic in the industry with semiconductors. That's why we had all these implementations out, and we wanted to make sure that the core piece was correct. We built test dies at AMD. We came up with -- there's not enough time to go into how it works, but we came up with a pretty elegant fix for that, and then we would channel that to semiconductors. One of the other companies in the collaboration didn't believe it was a problem, because they built one emulator and it didn't drift into sync, and that remained a problem until many, many, many, many years.

Paul Saffo

So, first of all, if there are audience questions and comments, actually, would everybody in this room who had something to do with the evolution of Ethernet raise your hand. Come on. There are some

people who haven't raised their hand that I know. That was the echoes of token ring. It's what happens when you do it in a ring. It's why trees are better.

Bill Hawe

This proves success has many fathers, but failure is an orphan. I think Ethernet was particularly promiscuous in that regard.

Paul Saffo

So, question, just a follow-on question for you, Bill, innovation, and Bob has written on this extensively, that we think of innovation, it's about technology, but of course it's about social processes, it's about interaction. Some of those can be quite colourful. We had a saying in the valley many years ago that for entrepreneurs, it's easier to ask for forgiveness than permission. Tell us about the time you got in a shouting match with Ken Olsen. I mean, I know there was more than one.

Bill Hawe

When one person is doing the talking and the other person is trembling, it's hard to call it a shouting match. What transpired was for some reason, the program manager thought it would be good for me to give a status report to the Executive Committee at Digital, at DEC. At the time, it was called DEC. And so he warned me that if Ken Olsen is humming or looking around, you're probably in trouble. And so I was giving this presentation on the status of what we were doing, and all of a sudden he was standing up at the end of the table, which seemed now to be as long as this room to me. It wasn't, but it sort of felt that way. And he says, "What does your badge say, son?" And I said, "Digital." And he says, "Does it say DEC or industry? Why are you giving away my technology to the industry?" And this was interesting, and the simulations I had were called the [DECWire] simulation. It wasn't called the Ethernet simulation, but nevertheless, we were trying to get it out into semiconductors and standardize it, but the standards were just really getting going. That's a whole other topic in IEEE.

Paul Saffo

I always thought that Ken's problem was he had a DEC wish.

Bill Hawe

I said, "Well, I think the strategy is it's better to get a big piece of a small pie than a small piece of a little pie, than all of a little pie." And so afterwards, I asked him, "Do you support this? Should I get on another project or something?" He said, "Oh, no, no, no. I just wanted to see if you were going to be a leader or anything, because you need to have arrows in the back from your friends, in the front from your enemies." It was great. Head games.

Paul Saffo

My gosh.

Radia Perlman

And just sort of a technical point, other than the radio stuff, the wi-fi stuff, mostly, Ethernet isn't CSMA/CD anymore. It's just point-to-point links with a half-hearted layer-3 protocol, which is this Spanning Tree. But that was because it wasn't designed to be doing that. But one of the real innovations of Ethernet was the six-byte address, which was a crazy idea, because IP only had a four-byte address and was intended to be the entire world, and Ethernet was designed for a few hundred nodes on a single wire. Why would they come up with six bytes? But this notion that you wouldn't have to configure things, that it would get born with a unique ID that you could carry around with you wherever you are was a really cool idea.

Unidentified Participant

Indeed.

Paul Saffo

And speaking to that, so about that radio thing, Dave, you really brought into this a radio aesthetic. As I recall, you're an Extra-class amateur radio operator, correct?

Dave Boggs

Yes.

Paul Saffo

And so to us on the outside, it looked like Ethernet, the idea was, "Oh, this is radio down coax." I know that's a grotesque oversimplification, but how did your radio perspective inform your design?

Dave Boggs

Well, it's not reliable, radio isn't, and somebody can always say, "Repeat, please. Something wiped it out." So retransmissions, and amateur radio people used to run message networks where they would accept written messages and send them by Morse code. I was quite familiar with checksums and retransmissions and all that, so when Bob said it to me, I said, "Oh, yeah, I understand all that."

Paul Saffo

It's fascinating to me that Bob came in, and he discovered ALOHAnet, but I think of Bob's many accomplishments, he's not a radio geek. Norm.

Norm Abramson

Well, Paul, if you go all the way back to Bob's initially rejected thesis at Harvard, you'll find a description of the architecture of what he was thinking of in his PhD thesis, and Bob used the term, which I subsequently adopted in a couple of papers -- it was never widely adopted, but I thought it was very insightful. He used the term "packet broadcasting," and so it seemed to me that right from the beginning, Bob was thinking radio but thinking radio on a wire. That is, he was broadcasting on the wire.

Dave Boggs

Bob talked about the ALOHAnet all the time. It was really the model he was trying to translate onto a wire.

Norm Abramson

He was grounded in radio, but he was constrained by the same thing that all of us were constrained about in those days, the FCC. The FCC would not allow you to do this, and it took some very inventive people within the structure of the FCC over the years to put us all on these little wireless units that we use today.

Paul Saffo

Fascinating. So we have a comment, so two thoughts I have is, one is the innovation process is the process of recombining things from other fields, really played out here. There was really a recombination of multiple perspectives. Also, the more I listen to this, the more I think about Bob's PhD getting rejected. I think I missed my opportunity. If I had gotten expelled from Harvard in freshman year, maybe I would have had a real career. Who knows. Sir, tell us who you are.

Bob Emmerson

I'm another Bob, Bob Emmerson, freelance writer. Earlier, a few minutes ago, Radia, you said that we don't have CSMA/CD -- I can't remember the acronym anymore. What happened to it? What allowed it to go away?

Radia Perlman

It just turned out that it's kind of evolved to just having intelligent switches and point-to-point links, just like how layer 3 was going to do. I was actually on a panel once at a conference. I was innocently there, and somebody said, "Oh, I want you to be representing Ethernet versus token ring." And so I read up -- at first, I said, "Oh, that's kind of hardware stuff and I'd rather not." They said, "Too bad. Everybody else has already said no, so you can't say no." I thought I could do a really good job of advocating for Ethernet against token ring, and then I found out there was a third guy on the panel, who was actually Sandy Fraser from Bell Labs, who I hadn't heard of at that point, and he was talking about star topologies, which didn't worry me. That was clearly stupid. It's a single point of failure. Then he talked first and just completely blew us away, so like with single point of failure, he

said, "Look at this topology. It looks like a single point of failure, but look at Ethernet, where every component is a single point of failure. Every transceiver, any one of them misbehaving, the whole thing is down." And so I was thinking, is he just sort of really smart and he could have taken any position and blown us away, or might he be right. So it turned out that he was kind of right, that the best way to cheaply hook up a building was with a more intelligent hub and hook everything together.

Unidentified Participant

Point-to-point links.

Radia Perlman

It's actually easier to isolate things.

Bill Hawe

Basically, all the things we argued against in the original thing, which is that picture that had a single wire and so forth and a twisted pair in a wiring closet. We were going to get rid of twisted pair and wiring closets. Whoops.

Radia Perlman

Right, and then if you couldn't have an intelligent thing with 10,000 ports, you'd hook them into each other to get more ports, and then you need the spanning tree, and so that's sort of how it evolved to be a better way of doing things. But the packet format remains, so it's not clear when people talk about Ethernet.

Bill Hawe

Actually, I have it here. This is the last remaining thing from the Digital, Intel, Xerox bluebook. It's page 27. It's the packet format. It's right here. It's the final thing that's left. However just one comment about that, though, it may sound pithy to say we were going to get rid of twisted pair and now we're on twisted pair, the Ethernet, the process and the core seeding, the core technology that seeded the industry and then the standards environment and a design which was just simple enough that pretty much anybody could get it to work, as opposed to you had to be a rocket scientist and everything had to be precisely the same, like a token ring Christmas tree lights in series kind of deal, that allowed an industry dynamic where Ethernet could evolve in a just-in-time environment. So when things run out of steam, there's a spring comes to the Ethernet and it reinvents itself. It gets reinvented. That's the engine that's going on now, and there isn't any particular reason for it to stop.

Paul Saffo

And that's a perfect final thought. The frustrating thing about today is I feel like we're just getting started on this panel, but, of course, we're out of time. So please join me in giving our panellists a hand.