Interview of William (Bill) Farr

Interviewed by:
James Pelkey

Recorded April 7, 1988
Boston, CA

CHM Reference number: X5671.2010

Bill Farr: I'm Bill Farr and I work for Prime Computer. I've been with Prime since the year of its inception, 1972, the year it opened its doors, and I've been involved with a number of product developments at Prime. Around 1975, I was sort of between projects, and Bill Poduska, who was Prime's VP of Engineering at that time, gave me some papers to read to try and seed, generate some ideas, seed my thinking. These papers were concerned with ALOHAnet in Hawaii.

James Pelkey: Abramson's papers?

Farr: The satellite network that they had developed to communicate with the mainland through radio, also Metcalfe's papers on Ethernet from MIT. I did a little more literature search and came up with some papers from Bell Labs on SpiderNet, and got really excited about the idea of actually linking computers together in a way that beyond just serial, RS-232 kinds of communication. We had thought about the ideas of linking computers with some high-speed parallel cable networks, where there was basically a master and slave. That was very limited, in that, in order to achieve any kind of performance, you had to have a lot of parallelism in terms of the data, and fairly high clock rates and expensive cables, and you were very limited in terms of the physical configuration of the network. The machines had to be close to each other. We actually developed, Paul Severino developed a product that we did sell. I've forgotten the name of it now, but it was limited to eight computers hooked up to a sixty wire cable, that was limited to, I believe, something like 100 feet of total cable length. We actually developed that product, tested it, and sold, I believe, 15 or 16 of them to customers. At the same time, I was working on these ideas that had developed from the ALOHAnet papers and Metcalfe's Ethernet papers and the token ring papers.

Pelkey: Do you remember what token ring papers? Was it the SpiderNet stuff?

Farr: It was the Spider stuff out of Bell Labs.

Pelkey: And Metcalfe's paper wasn't published in ACM until '76.

Farr: What I did was, I got his master's thesis from MIT.

Pelkey: How did you know to get that?

Farr: Bill Poduska pointed me. Bill was a full professor at MIT before he came to Prime. He was a full professor at MIT, and he was acting director of NASA's research center in Cambridge, and then he was head of -- when DOT took over the NASA work at Cambridge, he became director -- I believe he was director then - - of the Cambridge Research facility, and I had met Bill when I worked at Computer Controls, which was sold to Honeywell back in the early '60s.

Pelkey: So you worked with Art Carr and John Pugh?

Farr: Yes. A lot of us came from the CCC environment. I actually worked at DEC for a year back in the early '60s, too, so I knew Gordon Bell.

Pelkey: So Bill was aware of Bob's PhD dissertation, which was a refinement of the ALOHAnet concept --

Farr: Well, yes and no. It was an outgrowth of that thinking. That is, that one could -- the fundamental idea that ALOHAnet espoused was that one could put data on a medium at random, and then later on, see whether or not the data was corrupted, and then decide whether or not to retransmit. That was the fundamental idea, rather than trying to have some kind of gating protocol, some kind of master/slave relationship in a network that had nodes separated at great distances where it was very difficult to do that kind of synchronization, the idea was: "Let's just try it and see if it works, and if it works, fine, and if it doesn't work, we'll try it again," and you keep doing that until you succeed. If the network is lightly enough loaded, then that will be a successful strategy. I think that was the fundamental idea that ALOHAnet developed; the fact that it took seconds to get the data up to the satellite and back down, seconds is a long time when you're talking about computers and finding out whether things worked or not. There was no way to coordinate between the stations at either end when you do that kind of satellite transmission, so this idea of an unbounded network, the idea of transmitting data and then finding out later was not a totally new idea. That is, communication protocols had acknowledged messages that said whether or not data was successful end to end, but the idea of just allowing collisions, allowing corruption of data, was sort of a new idea. It was risky, I guess. Earlier, it had been thought that that was too inefficient to actually accomplish any significant high bandwidth transmission, but the ALOHA papers sort of indicated that this was a useful strategy in some ways, under some circumstances. Metcalfe's master's thesis --

Pelkey: Was it his master's or his PhD?

Farr: I believe it was his master's thesis, but I could be wrong. It sort of developed a very specific network, very similar to the Ethernet that is actually productized all over the world today, and he developed a number of statistical arguments in the paper that said that kind of a strategy would be successful, and actually, he built hardware, I believe, and made a prototype network work at that time. That was very interesting to me, but I've always grown up with the idea that you have to look to the future, and that over time, things will get faster, and you have to design things that will accommodate higher speeds, higher performance, and I'm very product oriented.

Pelkey: Excuse me, at that point in time, '75, were you aware of what was happening in Arpanet?

Farr: Yes.

Pelkey: And network control protocols? Did you familiarize yourself with those concepts and ideas and events at this time?

Farr: Not really in detail. I was familiar with Arpanet. Bill Poduska was very familiar with Arpanet. He had worked on Arpanet when he was at MIT. I guess I felt that Arpanet was telecommunications oriented, and I was looking for something that would be very much higher performance; a local area net, rather than telecommunications based. So I did a lot of thinking

about the pros and cons of an ALOHA type network or a slotted ALOHA type network, or Ethernet, or a token passing ring network, and my feeling was – and I did some analysis of the network loads; what would happen under peak loading conditions, under high data rates, conditions – that a token passing ring would be amenable to future growth, in terms of higher data rates and physical expansion, because the Ethernet protocols, the Ethernet implementation, having to wait after you had a collision, meant that the bigger the network got and the higher the data rates the network had, the less efficient Ethernet runs. That is, if you double your data rate on the network, you can't half the time that it takes to resolve a collision, so by doubling the data rate, you're running at half the efficiency on a heavily loaded network, so my view was that networking was for real and that, over time, the data rates would increase and the size of networks that would be required would increase, and a token passing ring scales in that environment. As you increase the data rate, the number of bits per second that you get on the network increases proportionally, and as you increase the size of the network, it doesn't affect the physical layer protocol, whereas with Ethernet, the bigger the network, the longer you have to wait to resolve collisions, and the higher the data rates, the more bits you destroy in that collision. So I felt that it was clear to me at the time that a token passing ring type of network was better for the future, so I wanted to try to build a product for Prime that would be based on that kind of a network, for those reasons. There were secondary considerations: I felt that the actual electrical circuitry needed to develop a token passing ring was simpler, in terms of you simply have to drive at one point and receive at one point, so you could control the environment on the cable. That is, you had a controlled driver and a controlled receiver, and you could pick whatever medium you wanted, whether it was microwave or cable or whatever point to point medium you wanted, fiber optics at that time was considered to be a -- beginning to be -- a viable alternative, whereas in an Ethernet environment, the driver driving to a basically undefined electrical characteristic, and the receiver having to discriminate between nodes that were very close and nodes that were very far away, and high signal levels and low signal levels, and multiple drivers and multiple receivers, just seemed to be a much more difficult task. Basically, I was working on this by myself with one other engineer, who was working on the driver and receiver circuitry. There were also some software people involved. So, to me, that seemed to be a much more realistic kind of network to develop for Prime at that time, plus as I say, the scalability was, to me, considered to be a real plus. So Bill encouraged further development, and in fact, I remember we had a discussion one day. He wanted to have a network to run at the National Computer Conference, the NCC show, which was some four or five months away, at that point.

Pelkey: Would that have been early '76?

Farr: I think that was late '75. It could have been '76, but I'm pretty sure it was '75, because I started working on this in late '74. I could be off by six months in there, I'm not sure, but I remember it was in June or July of whatever year it was. It was either '75 or '76 that NCC was scheduled. We had a heated discussion about that, whether that was a reasonable thing to do. I had done all the basic thinking about what the low-level network protocol was --

Pelkey: Had Bill and you interacted during this period of time so that he knew that you had, in fact, bitten the idea and thought that there was something there and had signed up to it conceptually?

Farr: Yeah, Prime was small enough at that time, so that I actually gave several presentations to the entire engineering community on the network proposal. I gave several presentations; one presentation on comparing the attributes of various kinds of networks and another on my specific proposal. So we had had several meetings by that time when the actual -- and I had written a specification at that point.

Pelkey: So it was time to get on it and get a product and let's show it!

Farr: Sure, there's not reason not to have it, and in fact, it could --

Pelkey: All the hard work had been done.

Farr: It could have been done, in terms of a basic prototype, because at that time at Prime, things got done that way. You just did it. There was no screwing around, but as I say, I'm very product oriented, and my goal was to get something that would be a robust product, not something that was a prototype, and my reasoning at this time was, if we develop this prototype, it would be interesting for the show, but then we'd basically have to go back and start over again and really do something real in order to have a shippable product. Anyway, it didn't get shown at NCC and, as a matter of fact, it took about a year from that point to actually have working hardware. One of the approaches that I took -- there was a discussion at the time as to whether, how to actually implement the low level physical protocol, whether to have phase-locked loops in each node and try to have the nodes synchronize themselves to the incoming signal, or something else -- and I chose basically an ASR-33 type of synchronization mechanism, where you have a sync character coming in, and when you get the sync character, you time out to the center of a bit cell, generate a strobe, a clock, and then you keep synchronous to that internal clock for a while, and then you resynchronize periodically. I didn't feel, again, with the resources that we had, that I could develop a stable phase-locked loop technology technique at the bit rates. That was another thing: I chose a bit rate of ten megahertz, which was really pushing the technology at that time. Bill Poduska would have preferred 20. I believe that was the number. Actually, he would have preferred 100, but he was willing to acknowledge that -- so one of the things that you are interested in is how things like this happen, how the ideas develop. Bill Poduska is one of the most creative people I've ever met. He was able to take an idea and push it way beyond what other people were thinking about, and see the value in very practical terms at the same time. He was always pushing. There was never a time where he was satisfied with anything. He would be satisfied with what you had done, but he always realized you could do more if you really wanted to. So there was this constant technological tension that he maintained, but it was a very fatherly type, fostering type. So he was always aware --

Pelkey: He must have been a real stimulus to the environment.

Farr: Yes he was. I also worked at DEC fairly early on, and Ken Olsen was very much the same type of person. As a matter of fact, Ken came around every day and talked to every engineer when I was there. Gordon Bell also was -- not quite as social, but very much the same kind of fostering person; very dynamic, generated lots of ideas and encouraged you to pick up on them and continue.

Pelkey: When did you have the working prototype done, do you recall?

Farr: It was sometime in 1976 that we had working boards, and it -- a very strange thing happened. As I said, I'm very product oriented and I could see the potential for this. I was very excited about the potential of actually being able to send data at very high rates between any number of computers within a fairly large facility and sharing that data. It struck me as something that everyone wanted, but no one could accomplish. Having RS-232 kinds of communications was so limiting, because it was strictly point to point, and not being able to share information between anyone who chose to look was extremely restrictive, so it seemed to me that this was an extremely exciting idea, and the fact that the technology had been developed to the point where this could actually now occur in a way that was comparable to what you could get when you access data from disk files, in terms of performance, seemed to me to be revolutionary, and I was very excited about Prime being one of the early developers of this technology, and productizing it. As I say, a number of companies had discussed this in the literature, but there was no way for the average computer customer to utilize this technology, so I was very excited. A number of us were, within the engineering community here, that this technology could actually be available to customers, and in fact, that we could use it internally, which was even more exciting, because then I could get the benefits of this.

Pelkey: So did you install it in here when you had this prototype?

Farr: Yeah. We originally developed -- let me go back in time. At the time we were developing this product, Xerox, in Rochester, NY, was one of our customers, and they heard about our network development. Now you have to remember that Xerox had Ethernet running in Palo Alto at this time, but Xerox, Rochester, NY, was talking to us about developing a large network for them. They saw it as a way of being able to communicate drawings, drafted drawings, storing drawings on disks throughout the company and being able to share these drawings and maintain them over a network. We could never figure out at the time, why they were interested in our network rather than Ethernet, but they seemed to feel that our network had more potential, in terms of performance, both in terms of the horsepower of the computing engines that we were selling versus what Xerox was selling at the time. Xerox was very much involved with personal -- precursors to workstations at that time, and they were basically small computer systems that lived in people's offices, with their own disks, the Altos systems. Prime was at the high end of the super minicomputer business, so Xerox -- I believe part of their motivation was the power of our processors, and they also believed, as I did, or seemed to, that the performance of the network would be better than the Ethernet product -- well Xerox didn't actually sell that as a product at that time. It did and it didn't. It was sort of a proprietary thing that you could get, but it was more hobby than a supported product, as I understood. So we were talking to Xerox and they were very interested in our network development, and we hadn't actually produced any hardware at that time. We were still designing the hardware, laying out the boards, and they were very concerned with whether or not this network would actually work, and whether or not you really could put 100 nodes or 200 nodes in a system. We didn't know for sure. All the working that I had done, all the thinking that I had done, said there should be no reason why it wasn't a scalable product, but strange things happen when you get into this unknown area. So we started out developing a product, and we build -- my idea was that if we

could get three nodes running, it would probably work for any number of nodes, but three nodes was sort of the critical mass. If you couldn't get three nodes running, then you didn't have a prayer, and once you got three nodes running, you had solved 99%, if not all, of the problems. But also, I was fairly sure that we would be surprised at some point. What we were looking at was, this kind of a network, a ring network, didn't seem to have any [limitations] in the curve. All previous networks, if you do any kind of analysis of networks, you get up to a certain throughput rate and the network just fails at some point and just won't support any more throughput. What we were looking at with these data rates was that this just wouldn't happen, that this was a scalable thing, because what you basically had was two computers communicating over a shared medium and the higher the traffic rate just meant that it might take little longer to get the data between any two nodes, but that you really wouldn't notice very much slowdown, if any, until the network got very heavily loaded, and at ten megabits per second, we just didn't see that that was going to be a problem. We were wondering whether three nodes would work, and the studies that I did said: "Gee, we should be able to get 50 or 100 nodes without seeing any performance degradation," and that was just an extremely exciting idea. So anyway, I was working on trying to get two nodes -- the first step was to get one node working, because the node had a built-in loopback capability, so I could actually test the node by itself, and then we took a giant step and went to two nodes over a cable, which was short cable at first, and we had a surprise there because we found that the token actually had to live on the ring. The system had to be big enough to support enough bits so the token would live and circulate continuously. We got that working, and then we went to three nodes, and we found a few more surprises in the low level protocol, and got that working --

Pelkey: It's now the end of '76?

Farr: Yeah, it was in there. I think it was the fall of '76 at this time.

Pelkey: You have three nodes working with a low level protocol.

Farr: And we had some software people working with us. There were three people: Dick Giles, Joe Brownstein and Paul Levine who had been working on X.25 around that time, and they were developing a communications protocol layer to interface to this local area network as well, so we were sort of working in parallel at that time. Once I got a two-node network running, you could actually develop the software in parallel to the hardware. Another very strange thing happened. It all of a sudden now became a big project, because we had to have a number of computers and a number of boards built just to test the rest of this network, so the resource requirement, the financial requirements, all of a sudden got beyond the scope of my small group. I had a boss at that time, Charley Smith, who -- I think he basically said: "Bill's been right often in the past. Maybe I should listen and support this," and we went to marketing, our marketing organization, and they saw absolutely zero need for this product. They thought it was a silly idea, and I think the basic reason for that was that nobody else had a network, so how could they sell something that they had no competition for.

Pelkey: Had Bill left by then?

Farr: No Bill was still here at that time. I said: "Can't you see that this is something that is valuable because people can share information at high speeds?" They said: "Well, people don't need to do that. Nobody's doing that now, why should -- why is that important?" It was very frustrating, because until we could get this targeted as a product, we couldn't get real support, even though we were working on it, about five people at that time within engineering, Prime was still a relatively small company, and we had very limited financial resources, and there were very few projects within Prime that weren't targeted for specific markets.

Pelkey: Now, the 400 was your big product?

Farr: That was about the same time.

Pelkey: So in '76, you introduced the 400?

Farr: Yes, I believe that's the case, or it may have been '75.

Pelkey: So you were starting to go up your ramp to $20 million in revenues around this point in time?

Farr: Yes, and we only had, I don't know how many systems in the field, but it was a relatively limited number, so we just didn't have a huge engineering budget, and we really needed, and I felt strongly that products should be developed as products, we couldn't afford experiments or failures or unsuccessful products, so we became very frustrated that we couldn't get support from the marketing organization for this, but a number of us felt strongly enough about it -- as a matter of fact, I think Hewlett Packard, about this time, was starting to say that they had a token ring network that they were productizing, but even that -- Hewlett Packard wasn't our main competitor. Our main competitors were Data General and DEC. We also heard that Data General was working on a network at that time, and a similar network, but since there was no product in the marketplace and there was no competition, there was no real pressure. So, I said to my boss Charley that we should buy 16 nodes and that we should produce a 16-node network. I wanted to go for 64, but I felt 16 was a reasonable thing to do, really outside of what we could do, so he signed up, anyway, for buying. He said: "I'll buy the boards. I'll buy the network, and we'll find the computers to put them in," because neither the manufacturing organization nor the marketing organization was willing to put up the money to build these boards. So we did. We produced 16 PrimeNet node controllers and we installed them in our engineering systems, and we had, of course, from the earliest days, we had developed timesharing computers, so every engineer had a terminal on his desk, and we installed these ring controllers in the computers, and now all of a sudden, all of these terminals were networked. The engineers loved it. They could send code and memos and all sorts of things back and forth and share all this data on all these disks. In other words, before, if you wanted something, you had to go pick up a disk pack and move it to another drive, and all of a sudden, you didn't have to do that, so within six or eight months --

Pelkey: What time frame is this?

Farr: This is early '77. We had all of a sudden become totally dependent on networks. It was like it happened overnight. It went from just an idea to: "How could we ever get along without this," and even then, the marketing organization thought this just wasn't worth anything. It became more and more clear during that year, that engineering was using this product to great advantage, and was totally dependent on it. We couldn't develop products without this anymore, especially in the software organization, where sharing of data was extremely important, and being able to look at other people's work, and audit things, and communicate rapidly a lot of detailed information. It became a staple of our engineering environment, and yet we couldn't get a product. We couldn't get marketing to sign up to sell this product. At that time, Hewlett Packard was producing great, glossy four and five page ads for this kind of thing. It could have been our ads. It was the same kind of product, and Data General had a product, and MITRE was talking about their network, MITRENET, which was a slightly different kind of product. Even Honeywell was producing a ring network at that point. I have forgotten what they called it, but they had an ad which said: "It's simpler. You just plug this connector into this connector and all of a sudden, your computers talk to each other." That was their thrust at that time. Even then, we felt very frustrated, because we felt we had a better product, that we were -- we weren't first in the marketplace, we were second, HP was first, but we had a great story to tell, and we got no support from out marketing organization, and it was very frustrating.

Pelkey: Now, where was Bill?

Farr: Bill was still here.

Pelkey: Who was president?

Farr: I believe at that time that Ken Fisher was president.

Pelkey: Why, given that Bill is the one that came to you originally and was a visionary and saw it was happening and then your engineering organization --

Farr: Bill was behind it. I don't want to go into this too much at this point, but there was a lot of -- one of the reasons that Bill left, I believe, was because he felt that he couldn't influence the company's direction significantly anymore. Bill was, at that time, beginning to think about workstations, and how personal workstations was an important thing in the future, and windowing -- being able to have multiple environments on a screen at that time --

Pelkey: Certainly your own environment here must have been -- plus, he must have known what was happening at Xerox PARC?

Farr: Oh, yeah.

Pelkey: The two of those must have very much shaped his thinking about what the next computer environment was going to look like.

Farr: Right, and he was in fact trying to do that kind of development within Prime, at that time, and not being very successful at getting away from the super mini --

Pelkey: But Prime was being so successful with the 400 and things were going great and we can be another DEC, a big minicomputer company --

Farr: Right, and why should we invest heavily in a new technology --

Pelkey: When we've got a winner on our hands?

Farr: So it's an understandable situation. Well, anyway, we did offer for sale this product. In fact, the IPC, the product that I told you that Paul Severino had developed, got pulled back at that time, because this was such a better product, we actually took back the 15 or 16 networks that we had supplied to customers, these tightly coupled networks, and gave them ring controller boards for their system, and updated their system software so that they had our new product rather than the one we had developed. Xerox did not, in fact, buy our network, I think partly because of internal pressure to use Ethernet, although it wasn't clear at all that Xerox was going to market Ethernet. Xerox had a history of producing outstanding products and never successfully marketing them, which helped the computer industry greatly, because other people saw the benefits of those ideas and developed product, and Xerox never seemed to be able to capitalize on their ideas in the computer industry. So, anyway, we did develop this product, we did sell it. We were basically selling it to people that wanted a few nodes. It was adopted very cautiously by our customers. Then the next big thing that happened was that Ford Motor Company started talking to us about networking. We had been courting Ford Motor Company for years. Our central region had targeted any of the automobile manufacturers as customers for us. They were fairly conservative and we just couldn't get our foot in the door. Well with our networking, Ford Motor Company said: "This looks like something that we can really use," and we started to talk seriously about large networks, and they also were concerned about how big a network could we actually support, because at this time, we were still talking about 20 nodes. We hadn't tested anything -- it's very hard to test a large network. Financially it's very hard, logistically it's very hard to do -- just what to do; how do you test such a thing? We have that problem even today. I think the largest PrimeNet network that exists is about 180 nodes at this point, still using the original controller boards that were shipped in 1977, '78, '79. We last year introduced a new version of this board that uses the same protocols, and they run compatibly, and it's microprocessor based and has more diagnostic capability, but it basically is the same kind of design. You can have the two boards mixed in the same system. So Ford Motor Company was interested, and they kept saying, as Xerox had: "Prove to us that this is going to work when you scale it up," and we said: "We can't because we've never done it, but there's no reason why it shouldn't, and we're committed to supporting this product." Anyway, they bought -- they started to buy our machines, and I believe --

Pelkey: As a consequence of your having networking? Did this make an impact on marketing at this point?

Farr: Yes, marketing was beginning to be believers. There were also other customers around the country: Citibank, Premier Systems and other customers, who wanted more computing power, but basically couldn't split up jobs in small enough pieces, so they had to rely on networking to share data and to effectively be able to use the MIPs available, not in a multi-

processing sense, but at least in terms of not breaking your tasks down into too small pieces, and being able to work on the same pieces of data. Ford had some ideas about developing a design network, where designers would send part information, or databases on parts, from designer to designer. A designer would work on a fender, or a brake drum, or something, and then take the data that he developed, the design database, and ship it to some other workstation, effectively, design C, then to perform other operations on it; actually develop molds or to analyze stresses or temperature, or develop economic models for producing this or whatever, and then eventually to send this database to their suppliers to make the part exactly as it was designed. And Ford actually did, over the years, carry that out and implements that today. In fact, today, their suppliers are required to have Prime computers with this network in order to qualify for certain parts of their product, to supply certain products, and I believe Ford is our largest customer today. I'm not sure of that, but they're certainly, if they're not our largest, they're our second largest, and they have a network that is as large as the one that we have internally, which, as I say, is today about 180 nodes and growing. I think they install a node a week in their network.

Pelkey: And the network stayed robust? What kinds of network software did you use? Now it's TCP/IP going to OSI.

Farr: We originally developed X.25 protocols. We were running a layered protocol. We had the physical layer, which was built into the board itself, and then we had a network interface module called a FAM, and then we had X.25 on top of that, so we were actually able to run local networks and area networks . . .

Tape side ends

Farr: I think probably, about three or four years ago, we realized that we would have to support Ethernet, we would have to support other networking protocols other than what we called PrimeNet, PrimeNet being X.25 compatible at some level, so we have, as a company, now embraced the idea of multi-vendor networks and multiple protocols and the layered system approach, the ISO approach, and we are actively working in that area to develop these standard interfaces. We also have our own internal interface layers, because we feel that, for performance reasons, you can't always go through every layer, so we provide the standard interfaces, but we will bypass layers where it's a reasonable thing to do within our own environment.

Pelkey: When did you develop electronic mail?

Farr: We had a mail product that we used internally, again -- I can't say exactly when it was, but I would say it was around 1980.

Pelkey: So it was some years after you got your network up?

Farr: Yes. We had transparent file transfer, but we had purchased what was called an office automation product at the time.

Pelkey: Do you recall from whom?

Farr: Yes, it was a company called ACS America in New York, just a small company, and we were very anxious to get into that business, and we cast about looking for a company or a product. There were a number of people in Prime who wanted to develop one internally, but at that time it was felt that that was not a high priority kind of thing, so we bought this product that turned out to be very poor, and we spent many years trying to make it work and patch it up. We have abandoned that product and have developed other products, but we had internal mail services fairly early, but we didn't offer them as products.

Pelkey: When you put a net up, I would imagine your engineers did a hack E-mail package.

Farr: Yes, very early on we did that, but there was no product.

Pelkey: Right, but I was thinking more internal than product. Were you aware of the E-mail concept when you created your network, or did that --

Farr: No, I was thinking more of transferring files and databases around the network, sharing data rather than actual individual workstation kinds of activities.

Pelkey: I would presume that E-mail -- once some engineers did a hack of it, all of a sudden everybody would have wanted to improve it and E-mail started to become indispensible.

Farr: Yes, in fact, that's where I said we became totally dependent on this network. That was one of the features that we had internally. The marketing would not even adopt -- we couldn't get them to use our mail system, even though we had computer networks there and here.

Pelkey: It's an incredible story.

Farr: I don't mean to knock the marketing organization severely, but it had such an impact on this product development that you can't help but be frustrated by it.

Pelkey: This is a common story. DEC, when they came out with their Blue Book in 1980, were you aware of that effort? You must have been aware in '79 when they announced they were going to do something, Xerox, Intel and DEC. Were you aware before the public pronouncements?

Farr: We were aware that there was something going on. It was kept fairly secret up until pretty close to the announcement. We knew about it maybe six months.

Pelkey: When you heard about that and saw who it was, what was your reaction?

Farr: By that time, we had a fair amount of momentum in selling our networks. It was not -- we were concerned that this consortium would, in fact, be a very powerful industry force, and that this might affect our customers, they might say: "We need an Ethernet, because that's the wave of the future," but at the same time, we realized that it wasn't necessary to have one physical layer network, that networking was something that was effective really at a much higher level. That, in order to be able to use networking, you had to have higher layer protocols, and

once you were able to produce those interfaces at these various levels, that the actual physical layer was not that significant. So we felt that we could produce multi-vendor networks with Ethernet, with any network that you wanted, at least through gateways, but at the same time we felt that the customers might perceive Ethernet as the wave of the future and be reluctant to invest heavily in ring networks.

Pelkey: Seeing that you had gone through the mental exercise and determined that rings were a better solution than Ethernet, and here your arch competitor picks a technology felt to be inferior.

Farr: That was certainly true. I felt, and I think other people felt, that ring networks did have advantages, but also by that time, we had been developing networking monitoring capabilities and we could measure peak and instantaneous traffic rates, and what we saw was that networks were really very lightly used, even when you have 50, 60, 80 computers on the network, the average data utilization was in the 10% range over a 24 hour period, and even the peaks were not very long. There is always a trade-off between peak data rate and how long you have to sustain that and how much capability you provide. Telephone networks had done that kind of thinking for 50 years. We saw that the peaks were only lasting for seconds and so that meant that the waiting times were very short, even during peaks. What we saw in the real world said that Ethernet would probably work just fine, so even though we said: "Yes, in the future, ring networks will probably be better than contention networks when the data rates go up, when the bit rates on the physical medium go up and as the networks get larger. We also felt that ring networks, although they had a fundamental Achilles heel, that is, any node in the network going down would take the network down, we had bypass capability in our networks from the beginning. You could bypass any node in the network, either under program control or manually. Also, Paul Levine and I developed a strategy whereby you could isolate any failed node in the network in order to bypass it. In other words, it doesn't do any good to be able to bypass unless you can tell which one to bypass. We knew that was doable and, in fact, had developed a strategy, although we hadn't implemented it in our released software. When we saw that Ethernet was having trouble in the actual implementation, that is, there was a problem early on, in terms of the saturation of the receivers. When you broadcast from a node and received it and an adjacent node on a large network, where you had to be able to receive this degraded signal at the end, there was this dynamic problem. We saw that there were a number of technical problems that had to be overcome if Ethernet was really going to be a successful product over the range of size of systems that we already knew were possible and were producing. So we knew that the Ethernet consortium was going to have some problems to resolve, not that they couldn't do it. We felt that they could and we felt that they would have a successful product if, in fact, they stuck with it, but we also knew that there was a history -- it was going to take time, and there was a history of Xerox not following through on things. In fact, there was a period of six or eight months when it looked like Ethernet might be a complete failure, that the industry was really not going to be willing to adopt it, and that there would not be a successful product. That was just about the time that Bill Poduska was starting up Apollo. Apollo was started with two things in mind -- three things: workstations, networked using commodity hardware and commodity technology, although they developed a proprietary operating system for their network, rather than using the commodity operating system. The commodity operating system idea was still three or four years away at that point. Apollo chose a ring network very similar to Prime's, with some enhancements. It's funny, because Bill Poduska had pushed for higher data

rates under the Prime network in 1975, and Apollo only achieved slightly marginal higher rates in the early '80s. I'm sure Bill was pushing just as hard for higher data rates at that time.

Pelkey: They were 12, right?

Farr: I think 12.5 actually, and we had been running with 10 since '76.

Pelkey: Do you think Prime's success in the marketplace caused DEC to be more aggressive in getting into the network business?

Farr: Yes I do. I think so. We were beginning to be noticed by DEC about that time.

Pelkey: This would be late '78, '79?

Farr: Yes.

Pelkey: And networking at this point in time -- you had networking out. When did Ford come aboard?

Farr: It was about the same time. Maybe 1980, because --

Pelkey: But there were big customers who were starting to pick Prime because, in fact, you had networking?

Farr: That's correct.

Pelkey: Were you aware that IBM was looking at token ring?

Farr: Yes, in fact, we knew that IBM had a token ring running in Europe, I believe in Switzerland in the '70s.

Pelkey: '74, '75 or later?

Farr: It's unclear. Both the Bell Labs and the IBM development of the token ring is very muddy, because they didn't publish it outside of their own companies. There were a few papers that were published, but there was very little in the literature.

Pelkey: Did you participate in any of the standards committee efforts?

Farr: No, we decided not to. Again, Prime was a fairly small company, and we didn't feel we could influence standards, and we didn't feel -- we felt that standards were something that, to a small company, were very inhibiting, that they would sap lots of energy and talent and would always end up being behind what we could accomplish ourselves. Standards are great, but they only adopt things that have already been successful.

Pelkey: In the early -- the IEEE LAN Committee -- there was meant to be one standard, and then there was a real big debate about Ethernet versus token ring -- token technology versus collision technology, and it had to be a two thirds vote, and while token got more than 50%, it didn't get the two thirds. There was an impasse, so they decided to switch the committee from being one standard to being multiple standards, and that's when IBM got into putting token ring out there also as a standard, because Ethernet was going to become a standard, and then token bus became a standard, and we ended up with these three technologies being standardized, rather than one kind of LAN technology being the standard of the IEEE.

Farr: That's correct. We were aware that that was going on. We didn't feel we could influence it because there was the Ethernet consortium and the IBM consortium, IBM being IBM, and we had a successful product and we were confident that it was the way to go. We weren't going to change until there was clear motivation to change, and we thought we'd wait until the dust settled before -- we figured that IBM had enough force to make a token ring product successful in the marketplace by themselves, and that's, in fact, I think, what happened. I wasn't party to the IEEE negotiations, but my understanding was that IBM came in and said: "Here's the standard," that they had developed, with full product specifications for their network, and they offered it to the committee with the idea that IBM would, in fact, adhere to this standard if the committee adopted it, and I think that was too much for the committee to resist. Here was a fully developed standard. Now, these standards committees always feel they have to contribute something --

Pelkey: But that was subsequent to the Blue Book?

Farr: Yes, but it was within a year --

Pelkey: Right, it was around that, the '81 time frame.

Farr: -- and I think that IBM's primary motivation, at that point, was to keep Ethernet from becoming a standard, because, I believe, they already had internal product plans for a ring network, and they didn't want the IEEE committee to adopt something else as a standard, so they were trying to railroad this, or at least keep the Ethernet standard from becoming the industry standard, because they had already committed to a product.

Pelkey: Were any people who were associated with the Prime LAN effort, did they ever spin off to other small companies, or become important in other companies?

Farr: Well, yes and no. Bill Poduska, of course, went and formed Apollo, which was a network-based workstation company. Paul Severino had been involved in networking, and then he went to work for a company called Data Translation, which really produced boards for PDP8s and PDP11s, mainly for analog to digital and digital to analog kinds of things, and then he went and formed his own company after that to develop local area network related products, the same kind of thing, developing networking products for DEC machines, and others. Microdata was one of his primary customers. A lot of the people that went to Apollo came from Prime and had been involved in the networking activity. Paul Levine is one. I think that was it. I can't think of any other companies that were network based, network product companies. One interesting thing that happened in, I think it was 1982: Bob Gordon, who worked here, he, I, and Paul

Levine had written a paper about Prime's Ringnet, which is our trademark network name, and he had also written a paper that he submitted to the system sciences symposium in Hawaii on the future of networking, and he couldn't go, so he asked me to go and present his paper, which I did. I actually presented a paper of my own as well as his, sort of combined, because it's hard to present somebody else's paper. There were people there from Xerox PARC. The theme of the conference was networking, and there were people there from Bell Labs and Xerox PARC --

Pelkey: What year was this?

Farr: '81 or '82. January, I believe, of '82. Basically, the paper that I presented was performance analysis of our network, and there was another fellow there who presented the performance analysis of the Ethernet from Xerox PARC. In fact, there was a paper published soon after that evaluating Ethernet's performance in various things, and the interesting thing about the performance measurements that these two papers presented was that the networks behaved very similarly. There was really no way to chose between them, and as I said, it's because most networking environments, even with large numbers of nodes, have very low average traffic rates. It was fascinating, because the characteristics of the networks were so different, and yet the results, in a real environment, were almost identical, in terms of up time, in terms of peak data rates, in terms of average data rates. We had a slightly larger network, so we were seeing higher data rates, but when you scaled it, it was very similar.

Pelkey: I have no other questions, Bill, you have been very helpful and I greatly appreciate your time. Unless there's something that you can think of that we haven't touched on?

Farr: No, I think that's pretty much it.

Pelkey: Thank you.

END OF THE INTERVIEW

